

Human-in-the-loop optimization of retinal prostheses encoders

Tristan Fauvel

▶ To cite this version:

Tristan Fauvel. Human-in-the-loop optimization of retinal prostheses encoders. Neuroscience. Sorbonne Université, 2021. English. NNT: 2021SORUS559. tel-03982094v2

HAL Id: tel-03982094 https://theses.hal.science/tel-03982094v2

Submitted on 10 Feb 2023 $\,$

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.







Human-in-the-loop optimization of retinal prostheses encoders.

Optimisation du codage du signal visuel dans les prothèses rétiniennes: une approche par apprentissage automatique.

Thèse de doctorat de Sorbonne Université

Présentée par

Tristan FAUVEL

Pour l'obtention du grade de Docteur en Neurosciences

Soutenue publiquement le 24 novembre 2021, devant un jury composé de :

Dr. Matthew CHALK, Institut de la Vision, INSERM	Directeur de thèse
Pr. Michael BEYELER, University of California, Santa Barbara	Rapporteur
Dr. Romy LORENZ, University of Cambridge	Rapportrice
Dr. Jean DAUNIZEAU, Institut du Cerveau, INSERM	Examinateur
Pr Ione FINE, University of Washington	Examinatrice
Dr. Peter NERI, École Normale Supérieure, CNRS	Examinateur
Pr. Günther ZECK, Technische Universität Wien	Examinateur

École doctorale Cerveau, Cognition et Comportement

Retinal degenerative diseases, such as retinitis pigmentosa or age-related macular degeneration, affect between 20 and 25 million people worldwide. These diseases lead to the gradual loss of photoreceptors, the light-sensitive cells of the retina, and therefore to blindness. Retinal prostheses are a promising strategy to restore sight to these patients. These devices are made of grids of electrodes or microphotodiodes positioned on or under the retina, or on the choroid -the vascular layer of the eye, to stimulate the remaining neurons of the retina by electrical impulses. The visual scene is filmed by a camera carried by the patient, and converted into an electrical stimulation pattern, to compensate for the loss of photoreceptors.

Despite promising beginnings and considerable technical progress, with the latest generations of implants made up of several thousand independent stimulation units, the visual performance of equipped patients remains well below expectations. Patients who no longer perceived light are now able to locate objects, perform visual recognition tasks or simple spatial navigation. However, the functional benefits remain very limited. Several reasons can explain this performance. First of all, the perception of shapes is greatly affected due to the diffusion of current in the tissue and the activation of the distal parts of the axons: a given electrode does not produce a 'pixel' in the visual field, but an elongated and ill-defined shape. In addition, the electrical stimulation of different types of retinal cells, which normally encode different information about the visual stimulus, is nonspecific, so downstream visual centers receive corrupted information. Extensive efforts have been made to obtain a more focused and specific stimulation, to process the incoming image to transmit only the information necessary for visual performance, and to attempt to mimic the neural code using an appropriate encoder.

In this thesis, we propose a new strategy for optimizing visual signal conversion in retinal prostheses based on the measurement of visual performance and patients' preferences. Users participate in a series of visual tasks, and their responses are used to continuously adjust the encoder according to a Bayesian optimization algorithm. Bayesian optimization is a powerful method to optimize functions whose analytical form is unknown without access to derivative information. It is especially used when the cost of a single function evaluation is high. It relies on a surrogate Bayesian model of the objective function which is used to query the system at locations informative about the optimum. The choice of querying a particular point is driven by a heuristic aiming at balancing exploration and exploitation.

In this thesis, we validate this strategy in participants with normal or corrected vision, using a prosthetic vision simulator. We show that preference-based optimization improves the quality of participants' perception and that this subjective improvement is transferred to stimuli other than those used during optimization, and is accompanied by a better visual acuity. The use of an adaptive sampling scheme allows faster optimization compared to random sampling. We used a parameterization of the encoder based on a model predicting the perception of patients equipped with an implant. We show that the optimization procedure is robust to errors in this model. This robustness, together with the fact that this method does not make any particular assumption regarding the type of implant, suggests that it could be implemented to improve sight restoration in patients. In addition, we show that an optimization strategy based on personal preference is more effective than optimization based on performance.

The challenges of applying preferential Bayesian optimization to retinal prostheses led us to develop new Bayesian optimization algorithms which outperform state-of-the-art methods in scenarios where the objective evaluation returns binary data, such as preference comparisons. In particular, many of the previously proposed method where either to computationally expensive to be used in a psychophysics context, or showed limited performance in practice. The new methods we proposed are based on the analytical decomposition of uncertainty about an evaluation outcome into its two components: aleatoric and epistemic, which allowed us to refine the definition of exploration in the context of Bayesian optimization.

The optimization of retinal prostheses encoders is an example of a situation where the optimized system can operate in many different environments, which induces several challenges for efficient and robust performance improvement. We explore this type of problem, in the case where the evaluation of the system involves binary measurements, by generalizing binary Bayesian optimization. We propose new heuristics combining methods from Bayesian optimization and active learning to efficiently optimize the objective across contexts. Les maladies de dégénérescence de la rétine, comme la rétinite pigmentaire ou la dégénérescence maculaire liée à l'âge, touchent entre 20 et 25 millions de personnes dans le monde. Ces maladies conduisent à une perte progressive des photorécepteurs, les cellules sensibles à la lumière de la rétine, et donc à la cécité. Les prothèses rétiniennes constituent une stratégie prometteuse pour rendre la vue à ces patients. Ces dispositifs sont des grilles d'électrodes ou de microphotodiodes positionnées sur ou sous la rétine, ou sur la choroïde, de façon à stimuler par des impulsions électriques les neurones restants de la rétine. La scène visuelle est filmée par une caméra portée par le patient, et convertie en un motif de stimulation électrique, de façon à compenser la perte des photorécepteurs.

Malgré des débuts encourageants et des progrès techniques considérables, avec les implants de dernière génération constitués de plusieurs milliers d'unités de stimulation indépendantes, les performances visuelles des patients équipés restent très en deçà des attentes. Des patients qui ne percevaient plus la lumière sont maintenant capables de localiser des objets, de réaliser des tâches de reconnaissance visuelle ou de navigation spatiale simples. Cependant, les bénéfices fonctionnels restent très limités. Plusieurs raisons peuvent expliquer ces performances. Tout d'abord, la perception des formes est très affectée du fait de la diffusion du courant dans le tissu et de l'activation des parties distales des axones: une électrode donnée ne produit pas un 'pixel' dans le champ visuel, mais une forme allongée et mal définie. De plus, la stimulation électrique des différents types de cellules de la rétine, qui en temps normal encodent des informations différentes sur le stimulus visuel, est non spécifique, de sorte que les centres visuels en aval reçoivent une information corrompue. De vastes efforts ont été entrepris pour obtenir une stimulation plus focalisée et plus spécifique, traiter l'image entrante de façon à ne transmettre que l'information nécessaire à la performance visuelle, et tenter de mimer le code neuronal à l'aide d'un encodeur approprié.

Dans cette thèse, nous proposons une nouvelle stratégie d'optimisation de la conversion du signal visuel dans les prothèses rétiniennes basée sur la mesure de la performance visuelle et de la préférence des patients. Les utilisateurs participent à une série de tâches visuelles, et leurs réponses sont utilisées pour ajuster l'encodeur en continu selon un algorithme d'optimisation bayésienne. L'optimisation bayésienne est une méthode puissante pour optimiser des fonctions dont la forme analytique est inconnue, et en l'absence d'information sur ses dérivées. Elle est particulièrement utile lorsque le coût d'évaluation de la fonction objectif est élevé. Elle repose sur un modèle bayésien de la fonction objectif qui est utilisé pour échantillonner la fonction à des emplacements informatifs quant à la valeur de l'optimum. Le choix d'un point particulier est guidé par une heuristique visant à trouver un compromis entre exploration et exploitation.

Dans cette thèse, nous évaluons cette stratégie chez des participants ayant une vision

normale ou corrigée, à l'aide d'un simulateur de vision prothétique. Nous montrons que l'optimisation basée sur la préférence permet d'améliorer la qualité de la perception des participants, et que cette amélioration subjective se transmet à d'autres stimuli que ceux utilisés lors de l'optimisation, et s'accompagne d'une meilleure acuité visuelle. L'utilisation d'un schéma d'échantillonnage adaptatif permet une optimisation plus rapide comparée à un échantillonnage aléatoire. De plus, nous montrons qu'une stratégie d'optimisation basée sur la préférence personnelle est plus efficace qu'une optimisation basée sur la performance. Notre travail fait usage d'une paramétrisation de l'encodeur basée sur un modèle prédisant la perception des patients équipés d'un implant. Nous montrons que la procédure d'optimisation est robuste à des erreurs dans ce modèle. Cette robustesse, jointe au fait que cette méthode ne fait pas d'hypothèse particulière quant au type d'implant, suggère qu'elle pourra être implémentée pour améliorer la restauration visuelle chez des patients.

Les défis de l'application de l'optimisation Bayésienne basée sur la préférence aux prothèses rétiniennes nous ont amenés à développer de nouveaux algorithmes d'optimisation bayésienne qui surpassent les méthodes existantes dans des scénarios où l'évaluation renvoie des données binaires, comme des comparaisons de préférence. En particulier, la plupart des méthodes existantes étaient soit très coûteuses en calcul pour être utilisées dans un contexte psychophysique, ou bien avaient de faibles performances en pratique. Les nouvelles méthodes que nous proposons reposent sur la décomposition analytique de l'incertitude sur l'issue d'une évaluation en ses deux composantes : aléatoire et épistémique, ce qui nous a permis d'affiner la définition de l'exploration dans le cadre de l'optimisation bayésienne.

L'optimisation de l'encodeur des prothèses rétiniennes est un exemple de situation où le système optimisé peut fonctionner dans de nombreux environnements différents, ce qui induit un certain nombre de difficultés pour obtenir une amélioration efficace et robuste des performances. Nous explorons ce type de problème, dans le cas où l'évaluation du système implique des mesures binaires, en généralisant l'algorithme d'optimisation bayésienne binaire. Cette généralisation s'appuie sur de nouvelles heuristiques combinant des méthodes issues de l'optimisation bayésienne et de l'apprentissage actif. First, I would like to thank my Ph.D. supervisor, Dr. Matthew Chalk, for giving me the opportunity to work on this fascinating project. Thanks for so much guidance and encouragement throughout these three challenging years.

I'd like to thank my team leader Olivier Marre, for timely feedbacks and insightful contributions on the main project of this thesis, as well as all my other teammates for providing such a rich scientific environment: Ulisse Ferrari, Elaine Orendorff, Baptiste Lefebvre, Danica Despotović, Francesco Trapani, Berat Semihcan Sermet, Thomas Buffet, Déborah Varro, Gabriel Mahuas, Gabriele Virgili, Takahiro Nemoto and Matías Goldin.

Special thanks to Romain Brette and his team for hosting me for most of my doctorate. I never thought I would learn so much about the neuroscience of paramecium by working on retinal prostheses.

I would like to thank Dr. Peter Neri and Dr. Alexandre Gramfort for their valuable advice as members of my thesis committee, as well as the members of my thesis jury: Pr. Ione Fine, Pr. Günther Zeck, Dr. Peter Neri, and in particular the two reviewers, Pr. Michael Beyeler, and Dr. Romy Lorenz.

Thanks to all the participants in my experiments, for their time and patience.

I would like to thank my friend Tristan Lazard who, I must admit, turned out to have a strong influence on my scientific (and athletic!) journey. For inextricable reasons, I would never have had the chance to work on this project if I had not met him some ten years ago.

I am very grateful to my parents for providing me with the best working environment I could dream of, twenty meters away from the seashore, when access to the lab was restricted because of the pandemic. Finally, I am very grateful to Sarah for her invaluable support over the course of my Ph.D. Sarah, you are the best belayer in life as in rock climbing!

My graduate studies were supported by a grant from the École Normale Supérieure.

 $x' \prec x$ is preferred over x'.

 $O(\cdot)$ The big-O asymptotic complexity of an algorithm.

$$[z]_+ = \max(0, z)$$

- θ Set of hyperparameters of a Gaussian process.
- *x* A *D*-dimensional input, i.e., $x \in \mathbb{R}^{D}$.
- y Vector of m observations, i.e., $y = [y_1, \dots, y_m]$.
- $\Gamma(z)$ The Gamma function.
- $\mathbb{E}(Z)$ Expectation of the random variable Z.
- $\mathbb{V}(Z)$ Variance of the random variable Z.
- \mathbf{A}^{-1} The inverse of a matrix.
- \mathbf{a}^{\top} , \mathbf{A}^{\top} Transpose of vector and matrix, respectively.
- $\mathbf{I}_{n \times n}$ *n*-by-*n* identity matrix.
- **K** Shorthand for the Gram matrix $K(\mathbf{X}, \mathbf{X})$.
- $\mathcal{B}(p)$ The Bernoulli law with parameter *p*.

 $\mathcal{GP}(\mu(\cdot), \Sigma(\cdot, \cdot))$ A Gaussian process with mean function μ and covariance function Σ .

- $\mathcal{U}(a, b)$ The uniform law on [a, b].
- \mathcal{X} Search space in the optimization problem.
- \mathcal{D}_t Data available at time *t*, made of input-observation pairs (*x*, *y*).
- Φ The cumulative distribution function of the unite Gaussian, $\Phi(z) = \int_{-\infty}^{z} \phi(t) dt$.
- ϕ The probability density function of the unite Gaussian, $\phi(t) = (2\pi)^{-1/2} \exp(-t^2/2)$.
- * The star is used to indicate the maximum of a function
- ** The double star is used to indicate that the maximum is restricted to the training set.

 $K(\mathbf{X}, \mathbf{X})$ The Gram matrix whose *i*, *j*th element is $k(\mathbf{x}_i, \mathbf{x}_j)$.

Notations

- $p(Z \mid W)$ Distribution of the random variable *Z* conditioned on *W*
- p(Z) Distribution of the random variable *Z*.
- T(h, a) Owen's T function.
- *y* An observation.
- tr(A) Trace of matrix A.

A	bstra	ct		ii
R	ésum	é		iv
A	cknov	wledge	ments	vi
Notations		vii		
1	Intr	oductio	on to retinal prosthetics	1
	1.1	Retina	al degeneration	1
		1.1.1	Structure of the normal retina	1
			The photoreceptors are the light-sensitive cells of the visual system \ldots	2
			The inner retina circuits shape the retinal output	2
			Ganglion cells transmit the processed image to the brain	4
			The Retinal Pigment Epithelium supports photoreceptors' function	5
		1.1.2	Photoreceptors degeneration	5
			Epidemiology of blindness	6
			Retinitis pigmentosa	6
			Age-related Macular Degeneration	8
			Remodeling of the degenerated retina	8
			Cortical plasticity following photoreceptors degeneration	10
		1.1.3	Approaches to vision restoration for retinal degenerative diseases	11
			Gene therapies	11
			Transplantation	12
			Optogenetic therapies	12
			Electronic visual prostheses	13
			Electrical stimulation of neurons	13
			Retinal prostheses	13
			Optic nerve prostheses	14
			LGN prostheses	14
			Cortical prostheses	15
			Cognitive prostheses and sensory substitution	16
	1.2	Visior	n restoration with retinal prostheses	16
		1.2.1	Prostheses designs	16
			Image capture	16
			Implant positioning	17

		Power and data transmission	17
	1.2.2	Prostheses function	18
		Controlling phosphenes	18
		Implant tuning	19
		Image translation	19
		Rehabilitation and perceptual learning	20
		Visual performance evaluation	21
		Visual outcomes for patients	22
		Visual outcomes with Argus II and IRIS II	22
		Visual outcomes with Alpha IMS and Alpha AMS	23
		Visual outcomes with PRIMA	24
		Visual outcomes with suprachoroidal implants	24
	1.2.3	Limiting factors for vision restoration	24
		Surgery	25
		Biocompatibility, charge injection limit and hermiticity	25
		Visual field	26
		Perceptual fading	26
		Head and eye movements	28
		Nonlinear mapping from RGCs eccentricity to visual field coordi-	
		nates	28
		Electrodes size and current spread	28
		Activation of passing axons	29
		Using the neural code	30
	1.2.4	Simulating prosthetic vision	33
1.3	Metho	ods for tuning visual implants	34
	1.3.1	State of the art of visual implants tuning	34
		Recordings-based strategies	34
		Approaches to the optimization problem	34
		From pattern optimization to encoder optimization	37
		Defining the target pattern	38
		Perception-based strategies	39
	1.3.2	State of the art in human-in-the-loop optimization	40
		A dialogue between human and machine	40
		Performance-based optimization	40
		Preference-based optimization	42
1.4	Thesis	soutline	43
Hur	nan-in-	the-loop optimization of visual prosthetic stimulation	16
2 1	Introd	Juction	4 6
4.1	211	Optimization of retinal prosthetic encoders	1 0 46
	2.1.1 212	Percention-based optimization of retina implant encoder	1 0 47
22	Frame	work	
2.2	221	Strategy assessment on normally sighted subjects	т, Д8
	4.4.1	oraces assessment or normany signed subjects	T 0

2

		2.2.2	Encoder parameterization	48
		2.2.3	Motivation of the closed-loop optimization approach	50
		2.2.4	Preferential Bayesian Optimization	51
	2.3	Metho	ods	52
		2.3.1	Principle of Preferential Bayesian optimization	52
		2.3.2	Maximally uncertain challenge acquisition	53
		2.3.3	Leveraging knowledge from previous patients.	54
			Constraints on the search space	54
			Kernel choice and hyperparameters selection	54
		2.3.4	Preference and visual acuity measurements	55
			Preference	55
			Visual acuity	55
		2.3.5	Bayesian analysis	56
	2.4	Result	ts	57
		2.4.1	Model-based encoding of visual inputs	57
		2.4.2	Preference-based optimization of retina implant encoder	58
		2.4.3	Transfer to other tasks	60
		2.4.4	Comparison with performance-based optimization	61
		2.4.5	Robustness to mis-parameterization of the encoder	63
	2.5	Discu	ssion	65
		2.5.1	Summary of contributions	65
		2.5.2	Related work	66
		2.5.3	Performance-based vs preference-based optimization	67
		2.5.4	Robustness of the adaptive preferential optimization procedure	67
		2.5.5	Limits and future developments	68
		2.5.6	Code availability	69
	2.6	Chapt	ter supplementary	70
		2.6.1	Instructions to the participants	70
		2.6.2	Supplementary figures	70
3	Prac	ctical B	ayesian optimization with binary and preference data	72
	3.1	Princi	ples of Bayesian optimization	73
		3.1.1	Standard Bayesian optimization	73
		3.1.2	Bayesian optimization with binary outputs	74
			Binary Bayesian Optimization	74
			Preferential Bayesian Optimization	75
	3.2	Gauss	sian processes as surrogate models	75
		3.2.1	Gaussian processes for regression	75
		3.2.2	Gaussian processes for classification	77
		3.2.3	Gaussian processes for preference learning	78
			Pairwise preference learning as classification	78
			Improving uncertainty estimation in Gaussian process preference learning	79
		3.2.4	Kernel and mean	82

3.3.1 Acquisition in standard Bayesian optimization 83 3.3.2 Exploration and exploitation in BO with binary outputs 84 Exploration and epistemic uncertainty 85 Exploration and epistemic uncertainty 85 Exploration and epistemic uncertainty 85 Exploration and epistemic uncertainty 87 3.3.3 Selecting the best point after spending the evaluation budget 89 3.3.4 Acquisition rules for binary BO 91 Binary Expected Improvement 91 Upper Credible Bound 92 3.3.5 Acquisition rules for PBO 94 Improvement-based policies 96 Bivariate Expected Improvement 97 Heuristics based on Thompson sampling 98 Duel Thompson Sampling 98 Maximally Uncertain Challenge 99 Comparing acquisition functions 100 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 104 3.5 Generalization to batches 105 3.6 Efficiently sampling from the maximum posterior distribution 108 <th>3.3</th> <th>Acqui</th> <th>sition rules with binary outputs</th>	3.3	Acqui	sition rules with binary outputs
3.3.2 Exploration and exploitation in BO with binary outputs 84 Exploration and epistemic uncertainty 85 Exploration and epistemic uncertainty 85 Epistemic uncertainty of a Gaussian process classification model 87 3.3.3 Selecting the best point after spending the evaluation budget 89 3.4 Acquisition rules for binary BO 91 Binary Expected Improvement 91 Upper Credible Bound 92 3.3.5 Acquisition rules for PBO 94 Improvement-based policies 96 Bivariate Expected Improvement 96 Copeland Expected Improvement 97 Heuristics based on Thompson sampling 98 KernelSelfSparring 99 Maximally Uncertain Challenge 99 Comparing acquisition functions 100 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 108 3.6.1 Efficiently sampling from Gaussian processes 108 3.6.1 Efficiently sampling from Gaussian processa spproximation		3.3.1	Acquisition in standard Bayesian optimization
Exploitation85Exploitation and epistemic uncertainty85Epistemic uncertainty of a Gaussian process classification model873.3.3Selecting the best point after spending the evaluation budget893.3.4Acquisition rules for binary BO91Binary Expected Improvement91Upper Credible Bound92Binary Knowledge Gradient923.3.5Acquisition rules for PBO94Improvement-based policies96Bivariate Expected Improvement97Heuristics based on Thompson sampling98Duel Thompson Sampling98KernelSelfSparring99Information-based policies99Maximally Uncertain Challenge99Comparing acquisition functions1003.4.1Binary Bayesian optimization1033.4.2Preferential Bayesian optimization1083.6.1Efficiently sampling from Gaussian processes108Weight-space approximate sampling109Sparse Spectrum Gaussian Process approximation109Sparse Spectrum Gaussian Process approximation110Acta201113.6.2Kernel Salproximation models1113.6.3Sompling from GP classification models110ActaAcquisition in adaptive contexts1184.1Introduction1184.1Introduction1184.1Introduction118		3.3.2	Exploration and exploitation in BO with binary outputs
Exploration and epistemic uncertainty85Epistemic uncertainty of a Gaussian process classification model873.3.3Selecting the best point after spending the evaluation budget893.3.4Acquisition rules for binary BO91Binary Expected Improvement91Upper Credible Bound92Binary Knowledge Gradient923.3.5Acquisition rules for PBO94Improvement-based policies96Bivariate Expected Improvement97Heuristics based on Thompson sampling98Duel Thompson Sampling98Maximally Uncertain Challenge99Anaming acquisition functions1023.4.1Binary Bayesian optimization1033.4.2Preferential Bayesian optimization1083.6.1Efficiently sampling from Gaussian processes108Meight-space approximate sampling1093.6.2Kernel approximation109Sparse Spectrum Gaussian Process approximation109Actical approximation109ActaSampling from GP classification models110ActaComparing approximate sampling1093.6.3Sampling from GP classification models110ActaSampling from GP classification models1103.6.4Comparing approximate sampling109Sparse Spectrum Gaussian Process approximation110ActaSoftware1113.6.4Comparing approximate sampling methods110ActaSomparing app			Exploitation
Epistemic uncertainty of a Gaussian process classification model873.3.3Selecting the best point after spending the evaluation budget893.3.4Acquisition rules for binary BO91Binary Expected Improvement91Upper Credible Bound92Binary Knowledge Gradient923.3.5Acquisition rules for PBO94Improvement-based policies96Bivariate Expected Improvement96Copeland Expected Improvement97Heuristics based on Thompson sampling98KernelSelfSparring99Information-based policies99Maximally Uncertain Challenge99Comparing acquisition functions1003.4.2Preferential Bayesian optimization1033.4.2Preferential Bayesian optimization1083.6.1Efficiently sampling from Gaussian processes108Weight-space approximate sampling1093.6.2Kernel approximation models1103.6.3Sampling from GP classification models1103.6.4Comparing approximation with non-Gaussian likelihoods110Accouled-bases approximate sampling1113.6.4Comparing approximate sampling methods1113.7Software1133.6.4Comparing approximation with non-Gaussian likelihoods110Acreal approximation on the preference kernels1113.6.4Comparing approximation with hon-Gaussian likelihoods110Acreal approximation with non-Gaussian likelihood			Exploration and epistemic uncertainty
3.3.3 Selecting the best point after spending the evaluation budget 89 3.3.4 Acquisition rules for binary EO 91 Binary Expected Improvement 91 Upper Credible Bound 92 Binary Knowledge Gradient 92 3.3.5 Acquisition rules for PBO 94 Improvement-based policies 96 Bivariate Expected Improvement 97 Heuristics based on Thompson sampling 98 Duel Thompson Sampling 98 KernelSelfSparring 99 Information-based policies 99 Maximally Uncertain Challenge 99 Comparing acquisition functions 100 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 104 3.6 Efficiently sampling from Gaussian processes 108 3.6.1 Efficiently sampling from Gaussian processes 108 Decoupled-bases approximate sampling 109 Sparse Spectrum Gaussian Process approximation 101 3.6.2 Kernel classification models 110 Meight-space auproximation with non-Gaussian likelihoods <td></td> <td></td> <td>Epistemic uncertainty of a Gaussian process classification model 87</td>			Epistemic uncertainty of a Gaussian process classification model 87
3.3.4 Acquisition rules for binary BO 91 Binary Expected Improvement 91 Upper Credible Bound 92 Binary Knowledge Gradient 92 3.3.5 Acquisition rules for PBO 94 Improvement-based policies 96 Bivariate Expected Improvement 96 Copeland Expected Improvement 97 Heuristics based on Thompson sampling 98 Duel Thompson Sampling 98 Comparing acquisition functions 90 Information-based policies 99 Comparing acquisition functions 100 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 103 3.4.1 Binary Bayesian optimization 108 3.6.1 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 Weight-space approximate sampling 109 109 S.6.2 Kernel Gaussian Process approximation 109 Sparse Spectrum Gaussian Process approximation 109 Sparse Spectrum Gaussian Process app		3.3.3	Selecting the best point after spending the evaluation budget 89
Binary Expected Improvement 91 Upper Credible Bound 92 Binary Knowledge Gradient 92 3.3.5 Acquisition rules for PBO 94 Improvement-based policies 96 Bivariate Expected Improvement 96 Copeland Expected Improvement 97 Heuristics based on Thompson sampling 98 Duel Thompson Sampling 98 Corpeland Expected Improvement 99 Information-based policies 99 Maximally Uncertain Challenge 99 Comparing acquisition functions 100 3.4 Results 102 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 108 3.6.1 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 Weight-space approximate sampling 109 109 Sparse Spectrum Gaussian Process approximation 109 Sparse Spectrum Gaussian Process approximation 109 <		3.3.4	Acquisition rules for binary BO
Upper Credible Bound 92 Binary Knowledge Gradient 92 3.3.5 Acquisition rules for PBO 94 Improvement-based policies 96 Bivariate Expected Improvement 96 Copeland Expected Improvement 97 Heuristics based on Thompson sampling 98 Duel Thompson Sampling 98 KernelSelfSparring 99 Information-based policies 99 Maximally Uncertain Challenge 99 Comparing acquisition functions 100 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 104 3.5 Generalization to batches 105 3.6 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 Decoupled-bases approximate sampling 109 Sparse Spectrum Gaussian Process approximation 109 Hilbert space Gaussian Process approximation 110 3.6.3 Sampling from CP classification models 110 A.6.4 Comparing approximation with non-Gaussian likelihoods </td <td></td> <td></td> <td>Binary Expected Improvement</td>			Binary Expected Improvement
Binary Knowledge Gradient 92 3.3.5 Acquisition rules for PBO 94 Improvement-based policies 96 Bivariate Expected Improvement 96 Copeland Expected Improvement 97 Heuristics based on Thompson sampling 98 Duel Thompson Sampling 98 KernelSelfSparring 99 Information-based policies 99 Maximally Uncertain Challenge 99 Comparing acquisition functions 100 3.4 Results 102 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 104 3.5 Generalization to batches 105 3.6.1 Efficiently sampling from Gaussian processes 108 3.6.1 Efficiently sampling from Gaussian processes 108 Decoupled-bases approximate sampling 109 Sparse Spectrum Gaussian Process approximation 109 Sparse Spectrum Gaussian Process approximation 110 3.6.3 Sampling from GP classification models 110			Upper Credible Bound
3.3.5 Acquisition rules for PBO 94 Improvement-based policies 96 Bivariate Expected Improvement 96 Copeland Expected Improvement 97 Heuristics based on Thompson sampling 98 Duel Thompson Sampling 98 KernelSelfSparring 99 Information-based policies 99 Maximally Uncertain Challenge 99 Comparing acquisition functions 100 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 104 3.5 Generalization to batches 105 3.6.1 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 Weight-space approximate sampling 109 109 3.6.2 Kernel approximation 109 Sparse Spectrum Gaussian Process approximation 110 3.6.3 Sampling from GP classification models 110 3.6.4 Comparing approximate sampling methods 110 3.6.3 Sampling from GP classification models 110 <			Binary Knowledge Gradient
Improvement-based policies96Bivariate Expected Improvement96Copeland Expected Improvement97Heuristics based on Thompson sampling98Duel Thompson Sampling98Super Standing99Information-based policies99Maximally Uncertain Challenge99Comparing acquisition functions1003.4Results1023.4.1Binary Bayesian optimization1033.4.2Preferential Bayesian optimization1043.5Generalization to batches1053.6Efficiently sampling from Gaussian processes1083.6.1Efficiently sampling from Gaussian processes1083.6.2Kernel approximate sampling1093.6.2Kernel approximation100Sparse Spectrum Gaussian Process approximation1103.6.3Sampling from GP classification models110Meight-space approximation with non-Gaussian likelihoods110Approximation of preference kernels1113.6.4Comparing approximate sampling methods1123.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts1184.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118		3.3.5	Acquisition rules for PBO
Bivariate Expected Improvement 96 Copeland Expected Improvement 97 Heuristics based on Thompson sampling 98 Duel Thompson Sampling 98 Maximally Uncertain Challenge 99 Maximally Uncertain Challenge 99 Comparing acquisition functions 100 3.4 Results 102 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 104 3.5 Generalization to batches 105 3.6 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 Weight-space approximate sampling 109 3.6.2 Kernel approximation 109 Scale Sparse Spectrum Gaussian Process approximation 109 Substimution of preference kernels 111 3.6.4 Comparing approximate sampling methods 111 3.6.3 Sampling from GP classification models 110 Meight-space approximate sampling methods 111 3.6.4 Comparing approximate sampling methods 111 <td></td> <td></td> <td>Improvement-based policies</td>			Improvement-based policies
Copeland Expected Improvement97Heuristics based on Thompson sampling98Duel Thompson Sampling98KernelSelfSparring99Information-based policies99Maximally Uncertain Challenge99Comparing acquisition functions1003.4Results1023.4.1Binary Bayesian optimization1033.4.2Preferential Bayesian optimization1043.5Generalization to batches1053.6Efficiently sampling from the maximum posterior distribution1083.6.1Efficiently sampling from Gaussian processes108Weight-space approximate sampling1093.6.2Kernel approximation109Sparse Spectrum Gaussian Process approximation1103.6.3Sampling from GP classification models110Meight-space approximate sampling1103.6.4Comparing approximate sampling methods1113.6.4Comparing approximate sampling methods1113.6.5Sampling from GP classification models1113.6.4Comparing approximate sampling methods1113.6.4Comparing approximate sampling methods1113.6.4Comparing approximate sampling methods1113.6.4Comparing approximate sampling methods1113.6.5Supplementary116Bayesian optimization in adaptive contexts1184.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs <t< td=""><td></td><td></td><td>Bivariate Expected Improvement</td></t<>			Bivariate Expected Improvement
Heuristics based on Thompson sampling 98 Duel Thompson Sampling 98 KernelSelfSparring 99 Information-based policies 99 Maximally Uncertain Challenge 99 Comparing acquisition functions 100 3.4 Results 102 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 104 3.5 Generalization to batches 105 3.6 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 Weight-space approximate sampling 109 Sparse Spectrum Gaussian Process approximation 109 Sparse Spectrum Gaussian Process approximation 110 3.6.3 Sampling from GP classification models 110 Weight-space approximate sampling methods 111 3.6.4 Comparing approximate sampling methods 111 3.6.3 Sampling from GP classification models 110 Meight-space approximate sampling methods 111 3.6.4 Comparing approximate sampling methods			Copeland Expected Improvement
Duel Thompson Sampling98KernelSelfSparring99Information-based policies99Maximally Uncertain Challenge99Comparing acquisition functions1003.4Results1023.4.1Binary Bayesian optimization1033.4.2Preferential Bayesian optimization1043.5Generalization to batches1053.6Efficiently sampling from the maximum posterior distribution1083.6.1Efficiently sampling from Gaussian processes108Weight-space approximate sampling1093.6.2Kernel approximation109Sparse Spectrum Gaussian Process approximation1003.6.3Sampling from GP classification models1103.6.3Sampling from GP classification models1103.6.4Comparing approximate sampling methods1113.6.5Sampling from GP classification models1103.6.3Sampling from GP classification models1113.6.4Comparing approximate sampling methods1113.6.3Sampling from GP classification models1113.6.4Comparing approximate sampling methods1113.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts1184.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118			Heuristics based on Thompson sampling
KernelSelfSparring99Information-based policies99Maximally Uncertain Challenge99Comparing acquisition functions1003.4Results1023.4.1Binary Bayesian optimization1033.4.2Preferential Bayesian optimization1043.5Generalization to batches1053.6Efficiently sampling from the maximum posterior distribution1083.6.1Efficiently sampling from Gaussian processes108Weight-space approximate sampling1093.6.2Kernel approximation109Sparse Spectrum Gaussian Process approximation109Sinarse Spectrum Gaussian Process approximation1103.6.3Sampling from GP classification models110Weight-space approximate sampling methods1113.6.4Comparing approximate sampling methods1113.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts1184.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118			Duel Thompson Sampling
Information-based policies99Maximally Uncertain Challenge99Comparing acquisition functions1003.4Results1023.4.1Binary Bayesian optimization1033.4.2Preferential Bayesian optimization1043.5Generalization to batches1053.6Efficiently sampling from the maximum posterior distribution1083.6.1Efficiently sampling from Gaussian processes108Weight-space approximate sampling1093.6.2Kernel approximation109Sparse Spectrum Gaussian Process approximation109Hilbert space Gaussian Process approximation1103.6.3Sampling from GP classification models110Weight-space approximate sampling1113.6.4Comparing approximate sampling methods1113.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts1184.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118			KernelSelfSparring
Maximally Uncertain Challenge99 Comparing acquisition functions1003.4Results1023.4.1Binary Bayesian optimization1033.4.2Preferential Bayesian optimization1043.5Generalization to batches1053.6Efficiently sampling from the maximum posterior distribution1083.6.1Efficiently sampling from Gaussian processes108Weight-space approximate sampling1093.6.2Kernel approximation109Sparse Spectrum Gaussian Process approximation109J.6.3Sampling from GP classification models1103.6.4Comparing approximate sampling methods1113.6.4Comparing approximate sampling methods1123.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts1184.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118			Information-based policies
Comparing acquisition functions1003.4Results1023.4.1Binary Bayesian optimization1033.4.2Preferential Bayesian optimization1043.5Generalization to batches1053.6Efficiently sampling from the maximum posterior distribution1083.6.1Efficiently sampling from Gaussian processes108Weight-space approximate sampling1093.6.2Kernel approximation109Sparse Spectrum Gaussian Process approximation109Sparse Spectrum Gaussian Process approximation1103.6.3Sampling from GP classification models110Weight-space approximation with non-Gaussian likelihoods1113.6.4Comparing approximate sampling methods1123.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts1184.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118			Maximally Uncertain Challenge
3.4 Results 102 3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 104 3.5 Generalization to batches 105 3.6 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 Weight-space approximate sampling 109 3.6.2 Kernel approximation 109 Sparse Spectrum Gaussian Process approximation 109 Software 110 3.6.3 Sampling from GP classification models 110 Weight-space approximate sampling methods 111 3.6.4 Comparing approximate sampling methods 112 3.7 Software 113 3.8 Conclusion 114 3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118			Comparing acquisition functions
3.4.1 Binary Bayesian optimization 103 3.4.2 Preferential Bayesian optimization 104 3.5 Generalization to batches 105 3.6 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 3.6.1 Efficiently sampling from Gaussian processes 108 Decoupled-bases approximate sampling 109 3.6.2 Kernel approximation 109 Sparse Spectrum Gaussian Process approximation 109 Sparse Spectrum Gaussian Process approximation 110 3.6.3 Sampling from GP classification models 110 Weight-space approximate sampling methods 110 Approximation of preference kernels 111 3.6.4 Comparing approximate sampling methods 112 3.7 Software 113 3.8 Conclusion 114 3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.11 Contextual Bayesian optimization with binary outputs 118	3.4	Result	ts
3.4.2 Preferential Bayesian optimization 104 3.5 Generalization to batches 105 3.6 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 3.6.1 Efficiently sampling from Gaussian processes 108 Weight-space approximate sampling 109 3.6.2 Kernel approximation 109 Sparse Spectrum Gaussian Process approximation 109 Hilbert space Gaussian Process approximation 110 3.6.3 Sampling from GP classification models 110 Meight-space approximate sampling methods 110 Approximation of preference kernels 111 3.6.4 Comparing approximate sampling methods 112 3.7 Software 113 3.8 Conclusion 114 3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118		3.4.1	Binary Bayesian optimization
3.5 Generalization to batches 105 3.6 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 Weight-space approximate sampling 108 Decoupled-bases approximate sampling 109 3.6.2 Kernel approximation 109 Sparse Spectrum Gaussian Process approximation 109 Hilbert space Gaussian Process approximation 110 3.6.3 Sampling from GP classification models 110 Weight-space approximate sampling methods 111 3.6.4 Comparing approximate sampling methods 112 3.7 Software 113 3.8 Conclusion 114 3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118		3.4.2	Preferential Bayesian optimization
3.6 Efficiently sampling from the maximum posterior distribution 108 3.6.1 Efficiently sampling from Gaussian processes 108 Weight-space approximate sampling 108 Decoupled-bases approximate sampling 109 3.6.2 Kernel approximation 109 Sparse Spectrum Gaussian Process approximation 109 Hilbert space Gaussian Process approximation 110 3.6.3 Sampling from GP classification models 110 3.6.4 Comparing approximate sampling methods 111 3.6.4 Comparing approximate sampling methods 112 3.7 Software 113 3.8 Conclusion 114 3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118	3.5	Gener	ralization to batches
3.6.1 Efficiently sampling from Gaussian processes 108 Weight-space approximate sampling 108 Decoupled-bases approximate sampling 109 3.6.2 Kernel approximation 109 Sparse Spectrum Gaussian Process approximation 109 Hilbert space Gaussian Process approximation 110 3.6.3 Sampling from GP classification models 110 Weight-space approximation with non-Gaussian likelihoods 110 Weight-space approximate sampling methods 111 3.6.4 Comparing approximate sampling methods 112 3.7 Software 113 3.8 Conclusion 114 3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118	3.6	Efficie	ently sampling from the maximum posterior distribution
Weight-space approximate sampling108Decoupled-bases approximate sampling1093.6.2Kernel approximation109Sparse Spectrum Gaussian Process approximation109Hilbert space Gaussian Process approximation1103.6.3Sampling from GP classification models110Weight-space approximation with non-Gaussian likelihoods110Approximation of preference kernels1113.6.4Comparing approximate sampling methods1123.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts1184.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118		3.6.1	Efficiently sampling from Gaussian processes
Decoupled-bases approximate sampling1093.6.2Kernel approximation109Sparse Spectrum Gaussian Process approximation109Hilbert space Gaussian Process approximation1103.6.3Sampling from GP classification models110Weight-space approximation with non-Gaussian likelihoods110Approximation of preference kernels1113.6.4Comparing approximate sampling methods1123.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts1184.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118			Weight-space approximate sampling
3.6.2 Kernel approximation 109 Sparse Spectrum Gaussian Process approximation 109 Hilbert space Gaussian Process approximation 110 3.6.3 Sampling from GP classification models 110 Weight-space approximation with non-Gaussian likelihoods 110 Approximation of preference kernels 111 3.6.4 Comparing approximate sampling methods 112 3.7 Software 113 3.8 Conclusion 114 3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118			Decoupled-bases approximate sampling
Sparse Spectrum Gaussian Process approximation109Hilbert space Gaussian Process approximation1103.6.3 Sampling from GP classification models110Weight-space approximation with non-Gaussian likelihoods110Approximation of preference kernels1113.6.4 Comparing approximate sampling methods1123.7 Software1133.8 Conclusion1143.9 Supplementary116Bayesian optimization in adaptive contexts1184.1 Introduction1184.1.1 Contextual Bayesian optimization with binary outputs118		3.6.2	Kernel approximation
Hilbert space Gaussian Process approximation1103.6.3Sampling from GP classification models110Weight-space approximation with non-Gaussian likelihoods110Approximation of preference kernels1113.6.4Comparing approximate sampling methods1123.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts4.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118			Sparse Spectrum Gaussian Process approximation
3.6.3 Sampling from GP classification models 110 Weight-space approximation with non-Gaussian likelihoods 110 Approximation of preference kernels 111 3.6.4 Comparing approximate sampling methods 112 3.7 Software 113 3.8 Conclusion 114 3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118			Hilbert space Gaussian Process approximation
Weight-space approximation with non-Gaussian likelihoods110Approximation of preference kernels1113.6.4Comparing approximate sampling methods1123.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts4.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118		3.6.3	Sampling from GP classification models
Approximation of preference kernels1113.6.4Comparing approximate sampling methods1123.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts4.1Introduction1184.1Contextual Bayesian optimization with binary outputs118			Weight-space approximation with non-Gaussian likelihoods
3.6.4 Comparing approximate sampling methods 112 3.7 Software 113 3.8 Conclusion 114 3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118			Approximation of preference kernels
3.7Software1133.8Conclusion1143.9Supplementary116Bayesian optimization in adaptive contexts4.1Introduction1184.1.1Contextual Bayesian optimization with binary outputs118		3.6.4	Comparing approximate sampling methods
3.8 Conclusion 114 3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118	3.7	Softw	are
3.9 Supplementary 116 Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118	3.8	Concl	usion
Bayesian optimization in adaptive contexts 118 4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118	3.9	Suppl	ementary
4.1 Introduction 118 4.1.1 Contextual Bayesian optimization with binary outputs 118	Bav	esian o	ptimization in adaptive contexts 118
4.1.1 Contextual Bayesian optimization with binary outputs	4.1	Introd	luction
		4.1.1	Contextual Bayesian optimization with binary outputs

4

		4.1.2 Problem formulation
	4.2	Acquisition rules for Bayesian optimization in adaptive contexts
		4.2.1 Knowledge gradient acquisition rule
		4.2.2 Sequential acquisition
		4.2.3 Generalization to preference judgements
	4.3	Experiments with synthetic test functions
		4.3.1 Binary Bayesian optimization in adaptive contexts
		4.3.2 Preferential Bayesian optimization in adaptive contexts
	4.4	Application: adaptive optimization in psychometric measurements
		4.4.1 Model of patients' responses 127
		4.4.2 Gaussian process model
		4.4.3 Results
	4.5	Discussion
		4.5.1 Summary of contributions
		4.5.2 Generalization to ConBO and Max-Min BO
		4.5.3 Related problems
		4.5.4 Code availability
	4.6	Chapter supplementary
		4.6.1 Supplementary results
		4.6.2 Kernel derivation for the psychophysics experiment
5	Die	135
0	51	Summary of contributions 135
	0.1	5.1.1 Validation of a human-in-the-loop optimization strategy for retina im-
		plant encoders in simulated prosthetic vision
		5.1.2 Uncertainty decomposition for efficient Bayesian optimization 137
		5.1.3 Contextual optimization and limitations of performance-based optimiza-
		tion
		5.1.4 Practical preferential Bayesian optimization
	5.2	Limitations of our approach
		5.2.1 Application to real-life patient
		5.2.2 Generalization to new generation implants
		5.2.3 Temporal coding
		5.2.4 Non-linearity of the perceptual model
	5.3	Leveraging domain knowledge
		5.3.1 Informed surrogate model
		5.3.2 Structured kernels
	5.4	Getting more from participants
		5.4.1 Ties inclusion
		5.4.2 From pairwise comparisons to top-k ranking
		5.4.3 Continuous tuning
		5.4.4 Everyday tuning
	5.5	Trends in Bayesian optimization

5.6 Conclusion	147
Appendix A Bayesian analysis of preference data	148
Appendix B Gaussian Conditionals	150
Appendix C Kernel definitions	151
Appendix D Gradient of the variance of the predictive class distribution for GP classi- fication	152
Appendix E Gradient of the knowledge gradient	153
Appendix F Conditional preference kernel F.1 Preference kernel F.2 Conditioning on the value function offset	155 155 155
Appendix G Finite-dimensional approximation to Gaussian processes G.1 Kernel approximation with random Fourier features	157 157 158 158
Appendix H Equivalence between the soft Copeland score distribution and the maxi- mum latent value distribution	160
Appendix I The Laplace approximation for Gaussian process classification I.1 Principle of the Laplace approximation I.2 Gaussian approximation of the posterior I.3 Approximate predictive distribution	161 161 162 163
Appendix J Benchmarks	164
List of Figures	167
List of Tables	168
Acronyms	169
Bibliography	199

1

Introduction to retinal prosthetics

Blindness is considered one of the most debilitating conditions. It has long been viewed as a fatality, and often depicted as a divine punishment (Rosse, 2010; McAllister, 2010; Jemison, 2016). However, with recent technological developments and the rise of artificial vision, curing blindness is no longer a dream. In the face of this tremendous challenge, progress in fields as diverse as neurophysiology, surgery, microelectronics, computer vision, or machine learning is converging to push the limits of sight restoration. Among the therapeutic strategies that are being developed, retinal prostheses are among the more mature and promising. These devices convert a visual signal filmed by a camera into electrical stimulation to replace lost photoreceptors. This thesis focuses on the problem of processing the visual signal down to a stimulation pattern via an encoder so that the elicited retinal activity conveys useful information to the patient. To address this problem, I developed and apply tools from the field of machine learning, in particular, Bayesian optimization. In the following sections, I will introduce aspects of the biology of the retina relevant to this thesis, as well as the main sight restoration strategies that are being investigated and the challenges they face. We will focus in particular on the state-ofthe-art in visual prosthetics and the problem of tuning the stimulation protocol. In chapter 2, I will present the results of experiments aiming at validating an automated encoder optimization strategy. In chapter 3, I will cover in more detail aspects of Bayesian optimization relevant to this work, in particular the limitations of previously used methods, and present developments that were crucial in the success of our encoder optimization method. A retina implant encoder should provide good prosthetic visual function across a wide range of tasks and contexts. Such contextual optimization comes with many challenges, in particular when contexts are not equally informative about the optimal configuration. In chapter 4, I will consider this type of problem, and demonstrate the efficiency of the solutions I propose on a problem from visual psychophysics.

1.1 Retinal degeneration

1.1.1 Structure of the normal retina

The human retina is a stratified tissue of about 200 µm thick that lines the back of the eye (figure 1.1). Cell bodies are arranged on three layers separated by two plexiform layers (made of dendrites and axons), the Outer Plexiform Layer (OPL) and the Inner Plexiform Layer (IPL). The outermost layer contains over 106 million photoreceptors (Curcio et al., 1990), light-sensitive cells in which light absorption initiates a signal transduction cascade. The Inner Nuclear Layer (INL) contains bipolar cells, which receive inputs from photoreceptors. The bipolar cells form synapses with the output cells of the retina, lying in the third nuclear layer, the Retina Ganglion Cells (RGC) (Curcio and Allen, 1990; Curcio et al., 1990; Nassi and Callaway, 2009). The axons of the RGCs converge towards the optic disc, where they form the optic nerve. From the retina, the visual signal flows down the optic nerves and through the chiasm, along the optic tract to the Lateral Geniculate Nucleus of the thalamus (LGN). From the LGN, signal then flows through the optic radiation to the Primary Visual Cortex (V1), and on to the higher visual areas. In addition to the bipolar cells, the inner nuclear layer contains horizontal cells and amacrine cells, interneurons which provide lateral connections in the OPL and the IPL respectively.

The photoreceptors are the light-sensitive cells of the visual system

The outer segment of photoreceptors contains light-sensitive proteins, opsins, which upon photon absorption initiate a transduction cascade that ultimately hyperpolarizes the cell and slows the release of neurotransmitter glutamate at the synapses with neighboring bipolar and horizontal cells. There are two types of photoreceptors: rods and cones, which are approximately 100 million and 6 million respectively in the human retina (Curcio et al., 1990). Rods have a long and cylindrical outer segment and are specialized for scotopic vision (i.e. in dim light) with single-photon detection ability (Rieke and Baylor, 1998). Cones have a shorter and conical outer segment and are only involved in photopic vision (i.e. in daylight). Humans have red or L-, green or M- and blue or S-sensitive cones with sensitivity spectra corresponding to long, medium, and short-wavelengths respectively (Baden et al., 2019). These distinct color specificities allow for color vision. The rods only express one type of opsin, so that they cannot differentiate from different spectra, which explains why color vision is not possible in dim light.

The photoreceptor spacing sets the sampling resolution of the visual system. The highest density of cones is found in the central region of the retina, called the fovea, and decreases towards the periphery (Curcio et al., 1990). The highest density of rods is found around the fovea and decreases both towards the periphery and the center of the fovea (Curcio et al., 1990). The density of bipolar cells and ganglion cells also decreases towards the periphery (Curcio and Allen, 1990), and their dendritic trees become larger (Rodieck et al., 1985).

The inner retina circuits shape the retinal output

The retina does not simply forward the photoreceptors' signal to the brain. Indeed, around 1.5 million ganglion cells are present in the human retina (Curcio and Allen, 1990). This implies that the input image is compressed between the photoreceptors layer and the RGCs layer. These computations are performed by the inner retina circuitry, and are responsible for spatial resolution, detection of color, motion and texture, light adaptation, control of circadian rhythms, and saccadic suppression (to counter image degradation by large eye movements) (Hood, 1998; Jasinska and Pyza, 2017; Gollisch and Meister, 2010). The photoreceptors to RGCs ratio depends on the eccentricity: the convergence is strong at the periphery, where a single RGC receives indirect inputs from hundreds of photoreceptors, whereas in the fovea, the ratio is closed to 1 (Watson, 2014). I will now briefly cover the function of the different cell types of



Figure 1.1: The human visual system (adapted from Airaghi Leccardi (2020). A.Structure of the human eye and neural organization of the human retina. C: cones, R: rods, EC: epithelial cells. B. Micrograph of a transverse section of a human retina (Sung and Chuang, 2010). IS: photoreceptors' inner segment, OS: photoreceptors' outer segment. Arrows indicate the path of the light. C. Schematics of the horizontal cut view of a human brain with the visual path highlighted (Mirochnik and Pezaris, 2019). ON: optic nerve, OC: optic chiasm, LGN: lateral geniculate nucleus

the inner retina.

Cone bipolar cells are of two main types, depending on their response to glutamate released by photoreceptors: OFF cone bipolar cells depolarize whereas ON cone bipolar cells hyperpolarize. This difference is due to the fact that OFF cone bipolar cells express AMPA and kainate receptors, which are cation channels opened by glutamate. ON cone bipolar cells, however, express the metabotropic glutamate receptor mGluR6. This receptor leads to the closing of the cation channel TRPM1 upon glutamate binding. Since photoreceptors are hyperpolarized by light, OFF cone bipolar cells are hyperpolarized by light, and ON cone bipolar cells are depolarized. There exist eleven types of cone bipolar cells based on their morphological and molecular properties (Breuninger et al., 2011). Importantly, the axon terminals of bipolar cells of different types terminate in different sublayers of the IPL, where they transfer the light signals onto the dendrites of amacrine and ganglion cells. Bipolar cells also receive inhibitory feedback from amacrine cells.

Rod bipolar cells express the metabotropic receptor mGluR6, as a consequence, they are depolarized by a light stimulus and are thus ON cells (Berntson and Taylor, 2000; Euler and Masland, 2000). Rod bipolar cells do not directly contact ganglion cells, but instead synapse with an AII amacrine cell (Famiglietti and Kolb, 1975; Raviola and Dacheux, 1987), which integrates inputs from many rod bipolar cells (figure 1.2). AII amacrine cells also form gap junctions with axon terminals of ON cone bipolar cells, and glycinergic synapses onto the axon terminals of OFF cone bipolar cells.

Horizontal cells, of which there are between one and three types in mammalian retinae, provide lateral interactions in the outer plexiform layer. Horizontal cells have a large dendritic arbor, they gather inputs from rods and cones and thus have receptive fields wider than their dendritic fields (Hombach et al., 2004). They are electrically coupled with other horizontal cells through gap junctions. Horizontal cells are hyperpolarized by light and form GABAergic synapses on photoreceptor synaptic terminals, providing inhibitory feedback proportional to

the mean luminance in their receptive field. The leading interpretation of this function is that it provides a mechanism of local gain control to the retina. They also contact bipolar cells, being responsible for the center-surround receptive field of bipolar cells (see 1.1.1).

Amacrine cells are inhibitory, usually axonless, interneurons. They receive inputs from bipolar and other amacrine cells and release neurotransmitters GABA and glycine. There are as many as 50 morphological types (MacNeil and Masland, 1998). Some types of amacrine cells are electrically coupled, allowing them to integrate signals from distant locations. They are involved in highly specialized computations. For example, amacrine cells are responsible for contextual effects for the responses of RGCs. This includes the classic center-surround antagonism, but also a variety of other effects (see Gollisch and Meister (2010) for a review), such as object motion detection by RGCs and direction-selectivity of RGCs (Euler et al., 2002). Many amacrine cells perform vertical integration between the layers of the IPL. They are for example responsible for crossover inhibition, the phenomenon by which they carry ON information into the OFF strata, and vice versa.

Ganglion cells transmit the processed image to the brain

The ganglion cells are located in the innermost nuclear layer of the retina, where they receive inputs from bipolar and amacrine cells. Their axons converge into a bundle, the optic nerve, that transmits action potentials conveying the information about the input image to the visual centers of the brain. OFF ganglion cells receive excitatory inputs from OFF cone bipolar cells in the IPL, whereas ON ganglion cells receive inputs from cone bipolar cells. This first functional distinction between ganglion cells dates back to Hartline (1938), who observed that some cells fire more spikes in response to light increments (ON cells) while others to light decrements (OFF cells). Note that there also exist ON-OFF cells which respond both to increments and decrements of light. The use of a localized stimulation (instead of full-field stimuli) such as a small light spot revealed the more complex spatial structure of RGCs receptive fields. RGCs receptive fields are circular and made of two concentric zones: the center and the surround, where effects of light are antagonistic (figure 1.2). What was previously classified as ON cells are ON-center, OFF-surround cells: the highest firing rate is triggered by a light increment in the center and a light decrement in the surround. This antagonistic center-surround receptive field is adapted to local contrast encoding.

RGCs types also differ depending on the temporal component of their receptive field. For example, the responses can be either transient or sustained (Cleland et al., 1971). The classification of RGCs types is still an ongoing research topic. For example, more than 30 different types have been reported in the mouse retina (Masland, 2012; Baden et al., 2016) which have been linked to morphologically and genetically defined cell populations. The best-characterized RGC populations include midget cells (with small, center-surround RFs), parasol cells (with large, center-surround achromatic RFs with high temporal sensitivity), and bistratified cells (with blue-yellow center-surround RFs), which collectively comprise around 90% of RGCs in the primate retina (Nassi and Callaway, 2009; Gollisch and Meister, 2010; Kling et al., 2020).

Other ganglion cell types include, for example, direction-selective ganglion cells (which respond to specific directions of motions), orientation-selective cells (which are selective to horizontal or vertical bars), or sparse intrinsically photosensitive cells (which are able to detect



Figure 1.2: Signal processing in the retina. A. RGCs' receptive fields organization (Kandel et al., 1991). B-C. Examples of retinal circuits (adapted from Wässle (2004), who modified it from Demb and Pugh (2002)). The ON rod pathways (B) and the OFF rod pathways (C). Red (resp. green) arrows indicate inhibitory (resp. excitatory) synapses.

light through the pigment melanopsin). These illustrate the diversity and complexity of the computations RGCs are involved in.

The Retinal Pigment Epithelium supports photoreceptors' function

The photoreceptor layer contacts the Retinal Pigment Epithelium (RPE), which is itself attached to the choroid (the eye's vascular layer). The RPE supports visual cells' function in many ways, the most important for our topic being the renewal of photoreceptors' outer segments. Indeed, phototransduction necessitates the reconstitution of the 11-cis retinal, the opsins' light-absorbing component, and this renewal process largely happens within RPE cells. The tips of photoreceptors are shed from their outer segments and phagocytized within the RPE (Young and Bok, 1969), in a continuous renewal process. The contents of the resulting phagolysosomes are incompletely degraded (Boulton et al., 1994) and the residuals form lipofuscin, which accumulates over life and may eventually lead to RPE cell death (Streeten, 1961).

1.1.2 Photoreceptors degeneration

The progressive death of photoreceptors leads to gradual loss of light perception, which may eventually lead to blindness (defined as having a visual acuity below 20/200).

Epidemiology of blindness

In two separate studies, the total number of people with vision impairment in 2010 was estimated at 191 million (Stevens et al., 2013) and 285 million (Pascolini and Mariotti, 2012) globally, with the number of those legally blind estimated at 32 and 39 million respectively. The causes of blindness vary widely from one country to another (Bourne et al., 2013). The most common cause of vision impairment among those aged 50 years and older in high-income regions in 2015 are uncorrected refractive error (prevalence of 3.03% (80% uncertainty interval (UI) 1.28%–5.3%), cataract (1.07%; 80% UI 0.31–2.53), Age-related Macular Degeneration (AMD) (0.81%; 80% UI 0.09–2.80), glaucoma (0.33%; 80% UI 0.04–1.15) and diabetic retinopathy (0.28%; 80% UI 0.02–1.10) (Bourne et al., 2018). AMD and Retinitis Pigmentosa (RP) are two of the most common retinal degenerative diseases (Hartong et al., 2006; Gehrs, 2010). They are both due to the gradual death of photoreceptors.

Retinitis pigmentosa

Retinitis pigmentosa (RP, see Hartong et al. (2006) for a review) denotes a group of genetic diseases that all lead to the progressive degeneration of rod photoreceptors followed by the gradual death of cone photoreceptors, while the inner nuclear layer degenerates later and partially (figure 1.3). The prevalence of RP was estimated to be 1:4000 in 2002 (Haim, 2002) and 1:3454 in 2013 (Bertelsen et al., 2014). The age at which patients report symptoms can range from early childhood to adulthood (Berson, 1993), and disease progression speed varies a lot, but most patients are legally blind by age 40 years (Grover et al., 1999). It is the most prevalent cause of visual disability in people younger than 60 years in industrialized countries (Pandova et al., 2019; Buch et al., 2004).

In RP, a mutation causes the death of rod photoreceptors. After rods are eliminated, cone photoreceptors slowly degenerate in a characteristic pattern, and after several years. Rod cell death causes night blindness, but visual disability results from cone degeneration (see Campochiaro and Mir (2018) for an excellent review on the mechanisms of cone degeneration in RP). Morphologic studies in human RP patients have shown moderate preservation of inner retinal cells and RGCs (Santos, 1997; Humayun et al., 1999).

Mutations that cause RP occur in a wide variety of genes. Many are involved in phototransduction, the visual cycle, photoreceptor structure, or gene transcription in photoreceptors (Wright et al., 2010). Mutations that cause RP lead to rod cell death, but not cone cell death, either because the mutated gene is differentially expressed in rods versus cones, or because it plays a more critical role in rods than cones. Mutations that affect both rods and cones lead to cone-rod dystrophies which have a phenotype distinct from that of RP (Berger et al., 2010).

The death of rods, which are metabolically very active cells, reduces oxygen consumption in the outer retina. This results in increased oxygen concentration in the outer retina (Yu et al., 2004; Padnick-Silver et al., 2006). The excess oxygen stimulates superoxide radical production by mismatches in the electron transport chain in mitochondria and by stimulation of NADPH oxidase complexes activity in the cytosol. Indeed, oxygen is the final electron acceptor of the electron transport chain in mitochondria. It receives electrons from four cytochrome c in the cytochrome c oxidase complex. With a high O_2 concentration, electron transfer to O_2 is more



Figure 1.3: Phenotype of retinitis pigmentosa (adapted from Hartong et al. (2006)). A. Fundus photograph of a normal retina. B. Fundus photograph of a retina with retinitis pigmentosa. Optic-disc pallor, attenuated retinal arterioles, and peripheral intraretinal pigment deposits are seen. C. Histological appearance of healthy human retina (left) and retina of a patient with mid-stage RP (right). The space between the retinal pigment epithelium and the outer nuclear layer in the diseased retina is a processing artifact.

likely to occur upstream in the chain, leading to the formation of superoxide O_2^- . The high levels of superoxide radicals overwhelm the antioxidant defense system and generate more reactive species. In particular, since neural tissue is rich in neurotransmitter NO, superoxide reacts with NO and generates peroxynitrite which is extremely damaging and difficult to detoxify (Pryor and Squadrito, 1995; Pacher et al., 2007). This results in progressive oxidative damage in both rods and cones which leads to their death (Komeima et al., 2006, 2007; Sanz et al., 2007).

Drugs or gene transfer that reduce oxidative stress promote cone survival and maintenance of function across several mouse models of RP with different mutations (Komeima et al., 2006, 2007; Oveson et al., 2011; Usui et al., 2009). Inversely, mutant RP mice deficient in superoxide dismutase (a component of the cell oxidative defense system) showed a more rapid evolution of the disease (Usui et al., 2011). Moreover, rods provide trophic support to other cells in the retina. In particular, rod-derived cone viability factors (RdCVF) are trophic proteins produced by rods, which promote cones' survival. Their absence makes cones more susceptible to oxidative (Elachouri et al., 2015) and metabolic stress due to deficient glucose uptake (Aït-Ali et al., 2015). Oral intake of antioxidative drugs such as N-acetylcysteine (NAC) is a potential therapeutic approach to reduce oxidative stress in cones, and hence slow down the disease progression (Lee et al., 2011).

Age-related Macular Degeneration

Age-related macular degeneration (AMD, see de Jong (2006) and Fleckenstein et al. (2021) for review) is the leading cause of severe vision loss in individuals above 55 years of age in the developed world. The age at onset of clinically apparent AMD varies greatly but typically begins at above 55 years of age (Ferris et al., 2013). A 2014 meta-analysis of studies of persons between 45 and 85 years of age estimated a prevalence of age-related macular degeneration of 8.69 % (95% credible intervals (CrI) 4.26–17.40%), with a prevalence of late-stage AMD of 0.37% (95% CrI 0.18–0.77%). Due to global population growth and aging, the projected number of people with AMD is expected to rise to about 288 million by 2040 (Wong et al., 2014).

The macula is a 5.5 mm diameter retina region in the axis of the pupil, which contains the highest concentration of cones. In AMD, the macula undergoes a considerable decrease in the number of photoreceptors, with less than 20% remaining (Medeiros and Curcio, 2001), leading to central vision loss (figure 1.4). This death of photoreceptors is paralleled by a degeneration of the RPE.

Two forms of AMD exist. Dry AMD (also called non-neovascular AMD, or geographic atrophy (GA), the most common, 80%) is characterized by the thinning of the macula and the accumulation of proteins in so-called drusen. In addition to drusen, a more diffused form of deposit known as subretinal drusenoid deposit is commonly observed in the subretinal space. Neovascular Age-related Macular Degeneration (NVAMD) is caused by abnormal angiogenesis (growth of vascular tissue) in the outer retina, subretinal space, or sub-RPE space. The exudative stage of NVAMD (also known as "wet AMD") begins when these new vessels leak or rupture, resulting in fluid accumulation and/or hemorrhages. Vision loss is faster with wet AMD. Both forms of AMD can coexist in the same patient and can evolve into the other form.

The mechanisms responsible for the onset of AMD are not well known, but it is believed that the slow accumulation of debris such as lipofuscin within RPE cells and in the extracellular space in the form of drusen trigger RPE cells death and local inflammatory reaction. Indeed, the parafoveal ring, where rod density is highest (Curcio et al., 1993) and where dry age-related macular degeneration often begins, has the highest concentration of lipofuscin in the retina.

There is no treatment for the dry form of AMD. In wet AMD, anti-vascular endothelial growth factor drugs can be injected into the eye to prevent angiogenesis. If the disease is detected early enough, laser surgery can also be used to destroy abnormal blood vessels. AMD is a multifactorial disease, risk factors include a diet high in saturated fatty acids, overweight, smoking, age over 50 years old, hypertension, and genetic factors (Seddon, 2017).

Morphometric analysis showed a 50% loss of ganglion cells in NVAMD patients, but no significant difference from normal subjects in non-NVAMD cases (Merabet et al., 2005), suggesting that vision restoration strategies based on the activation of ganglion cells could be adapted to this disease.

Remodeling of the degenerated retina

Accompanying and following photoreceptors degeneration, extensive remodeling occurs in the retina (reviewed in the case of RP by Marc et al. (2003) and Strettoi (2015), and in general by Jones and Marc (2005). For a recent review focused on late-stage degeneration, see Pfeiffer et al.

1.1. Retinal degeneration



Figure 1.4: Stages of AMD (adapted from (Fleckenstein et al., 2021)). A. Normal retina. The rectangular area corresponds to the macula. B. Macula with early stage AMD. Drusen are visible on the fundus (yellowish deposits, arrowheads). C. Geographic atrophy (arrowheads on the fundus). D. Neovascular AMD. Haemorrhage is visible on the fundus (asterisks).

(2020)). The term remodeling refers to the process of morphological and functional changes involving neurons, glia, and blood vessels in the retina that accompany and follow photoreceptors degeneration. Understanding this phenomenon is crucial for any sight restoration strategy aiming at compensating for photoreceptors loss. A retina implant, for example, must be able to accomodate ongoing changes that occur following blindness. Importantly, despite the diversity of causes underlying degeneration, the process of remodeling occurs in similar sequences of events, including in particular extensive rewiring of the inner retinal network. Concomitantly to the disappearance of photoreceptors, resident microglia is activated, engulf and clear dying photoreceptors. Müller cells (the most common type of glial cells found in the retina) fill the partially vacated space, as well as the subretinal space, forming a dense seal between the remaining retinal neurons and the RPE. These newly formed glial surfaces serve as substrates for neuronal migration and process extension. Surviving cones undergo sprouting of their telodendria, with some of the processes extending as far as the Ganglion Cell Layer (GCL), and might form ectopic synapses with rod bipolar cells. Axonal arbors of horizontal cells undergo either sprouting or regression depending on the models, as well as dendrite retraction. Bipolar cells display morphological changes, including dendrite reduction, retraction, and mislocation, as well as displacement of cell bodies. Their molecular phenotype also evolves, with decreased expression of mGluR6 receptors correlated with a decrease in glutamate-activated currents. Amacrine cells display neurite sprouting and form new, sometimes aberrant, synaptic connections. Some amacrine cells also migrate down into the GCL. Ganglion cells display limited morphological modifications, retain their intrinsic electrical properties, and projection to brain targets.

These remodeling events may be caused in particular by deafferentiation of horizontal and bipolar cells of photoreceptor inputs and local inflammatory reactions. They lead to abnormal activity of the retina. In particular, in the absence of photoreceptors, the retina is very active, indicating that it is self-signaling without light-evoked signals. Moreover, RGCs spontaneous activity shifts from a random pattern to a rhythmic one in which bursts of spikes occur at roughly 10 Hz (Goo et al., 2011; Menzler and Zeck, 2011).

Cortical plasticity following photoreceptors degeneration

Cortical plasticity (also known as neuroplasticity) describes the ability of the cortex to change its structure or function in response to experience. Cortical plasticity can be observed at multiple temporal scales, ranging from short-term (seconds to minutes) to long-term (days to many months) (Horton et al., 2017). Understanding the plasticity of the visual cortex in patients with retinal degenerative diseases is crucial to develop therapies based on the replacement of the photoreceptors inputs (see Beyeler et al. (2017b) for review). Indeed, remodeling of the visual cortex following sensory deprivation may preclude restoring functional vision, whereas cortical plasticity may allow the visual system to make sense of new, unnatural stimuli.

Brain imaging and transcranial magnetic stimulation studies have shown that after prolonged blindness, the visual cortex can be repurposed to other functions. It is activated in response to stimuli from other sensory modalities, either auditory (Weeks et al., 2000; Gougoux et al., 2005; Saenz et al., 2008) or tactile (Sadato et al., 1996; Büchel et al., 1998). Even for late-blind patients, the visual cortex responds to both auditory (Hölig et al., 2014) and tactile (Cunningham et al., 2011) stimuli. These observations raise the question of whether such cross-modal plasticity can impede vision restoration.

Case studies on patients who had their sight "restored" after prolonged vision loss by procedures such as cataract removal (Fine and Jacobs, 2002; Ostrovsky et al., 2009; Sinha et al., 2013; McKyton et al., 2015) or corneal replacement (Fine et al., 2003; Šikl et al., 2013) can give us some insights into the rehabilitation capabilities following prolonged blindness. Impairment in shape processing, object recognition, face processing, and space perception are observed in sight-recovery patients (Held et al., 2011; Šikl et al., 2013; Fine et al., 2003), and these persist more than a decade after surgery (Huber et al., 2015). Interestingly these deficits are observed both in individuals who had their sight restored in late childhood (Sinha et al., 2013) and adulthood (Šikl et al., 2013), and in the latter case both in individuals who had congenital or early blindness and who lost their sight in adulthood (Šikl et al., 2013). However, in some of these patients, improvements in visual processing followed in the months following surgery. These improvements were stronger for high-level visual tasks. About low-level visual function, improvements in the contrast sensitivity function has been noted in several individuals who had sight restored at a young age (between the ages of 8–17 years) (Kalia et al., 2014), but these improvements may be age-dependent as they were not observed in patients in their 40s (Fine and Jacobs, 2002; Fine et al., 2003).

The aforementioned studies were performed in patients with prolonged vision loss, and the cortex of patients with recent visual impairment may not display such an extended repurposing. Neurophysiological (Rosa et al., 1995; Horton and Hocking, 1998; Smirnakis et al., 2005; Masuda et al., 2008; Baseler et al., 2011) and fMRI (Rosa et al., 1995; Horton and Hocking, 1998; Smirnakis et al., 2005; Masuda et al., 2008, 2010; Baseler et al., 2011; Shao et al., 2013) studies have failed to find retinotopic reorganization within cortex following retinal lesions or retinal dystrophy. These results suggest that early intervention may allow sight restoration before the occurrence of visual cortex repurposing.

Another important question to address is whether cortical plasticity can facilitate vision restoration. Previous experience with cochlear implants showed that cortical plasticity is capable of compensating for significant sensory input distortion. Adult-implanted cochlear implant users initially report extremely unnatural and incomprehensible perceptual experiences. However, over hours to approximately a year, they adapt to this new, distorted sensory input (Hallberg and Ringdahl, 2004) and become better at tasks such as speech recognition (Shannon, 2012). These changes are attributed to cortical plasticity.

The visual system is usually considered less plastic than the auditory system (Beyeler et al., 2017b). The question remains whether the cortex of patients suffering from photoreceptors degeneration is plastic enough to learn to interpret the signal sent by vision restoration technologies. Cortical plasticity following photoreceptors degeneration has been investigated in both RP (Lunghi et al., 2019) and AMD (Maniglia et al., 2018) patients. In the first study, cortical plasticity was evaluated by observing the effect of short-term (2 hours) monocular deprivation on sensory ocular dominance, measured with binocular rivalry. In the second one, it was measured using lateral masking, a contrast sensitivity modulation induced by collinear flankers. In RP patients, the cortex showed plasticity in the normal range, and stronger visual impairment was associated with higher plasticity (Lunghi et al., 2019). AMD patients showed higher cortical plasticity compared to controls (Maniglia et al., 2018). Taken together, these observations give hope that cortical plasticity can facilitate the restoration of at least basic visual function in patients with recent blindness.

1.1.3 Approaches to vision restoration for retinal degenerative diseases

In the following section, I will briefly describe the various approaches to restore vision to patients suffering from retinal degenerative diseases. Other approaches aiming at preventing or slowing degeneration also exist and show promising results (see sections 1.1.2), such as interventions in biochemical pathways with neuroprotective agents (Sieving et al., 2006; Kauper et al., 2012).

Gene therapies

Gene therapies (reviewed in McClements and MacLaren (2013) and Petrs-Silva and Linden (2013)) are a promising strategy to cure inherited retinal degenerations such as RP or Usher syndrome. Indeed, in most cases these diseases result from pathogenic mutations in single genes leading to loss of protein function; it is thus possible to repair, replace, or compensate for the mutated gene to restore function (Gupta and Huckfeldt, 2017). In dominant diseases, in which the mutations induce a pathogenic protein, more advanced strategies such as gene silencing and genome editing strategies are required.

Eye diseases are an ideal target for gene therapy, compared for example to systemic diseases. Indeed, the vectors can be injected directly into the eye, either in the vitreus or subretinally. The eye being an immune-privileged organ (it can tolerate the introduction of antigens without triggering an inflammatory immune response), the risk of deleterious immune reaction to the vectors or newly expressed protein is mitigated (Anand et al., 2002; Bennett, 2003). Moreover, the injection can be performed unilaterally, limiting the risk for patients, and making use of the contralateral eye as a control.

This approach has been successful in the preservation or partial restoration of vision in Leber's congenital amaurosis (Kumaran et al., 2017) and choroideremia (Xue and MacLaren, 2018). In 2017, the voretigene neparvovec (Spark Therapeutics), designed to provide a functional copy of the gene RPE65, became the first gene therapy to complete a phase III clinical trial. One challenge that gene therapies face, however, is the huge number of identified genes causing inherited retinopathies.

Transplantation

A strategy to delay degeneration and eventually regenerate damaged photoreceptors is to replace the damaged RPE cells (see da Cruz et al. (2007) for a review). The strategy was assessed by transplanting intact sheets of the retina with RPE in a rat model of photoreceptors degeneration, leading to improved pupillary light reflex (Whiteley et al., 1996) and restored visually evoked cortical responses (Woch et al., 2001). It was shown that the transplanted RPE was able to support photoreceptor cells survival and initiate photoreceptors outer segments regeneration (Lin et al., 1996). One possibility is to transplant embryonic stem cell-derived RPE cells (Schwartz et al., 2012). This strategy has shown the ability to reverse visual loss in a limited number of cases by restoring photoreceptors' function, but without demonstrated visual benefits. Moreover, this strategy cannot restore lost photoreceptors. Another path to vision restoration is to directly insert precursors of photoreceptors from embryonic stem cells subretinally (MacLaren et al., 2006) (see Ramsden et al. (2013); Barber et al. (2013) and Pearson (2014) for reviews), or photoroceptors from a donor (Santos-Ferreira et al., 2015), with promising results in preclinical studies where mice recovered photopic responses following transplantation.

Optogenetic therapies

The optogenetic approach (reviewed by Duebel et al. (2015) and Fine et al. (2015)) aims at replacing the lost photoreceptors by making the remaining inner retinal cells photosensitive. This is achieved by making these cells express light-gated channels such as channelrhodopsin-2 or halorhodopsin using viral vectors (Degenaar et al., 2009; Busskamp et al., 2012). However, the light sensitivity of the transfected neurons is such that they require luminance well beyond natural light (about 100mW/cm²) for action potential initiation. In practice, the patient wears goggles carrying a camera and a projector that sends a light beam to the eye. In a clinical trial (Sahel et al., 2021), a RP patient received an intravitreal injection of an AAV vector encoding the optogenetic protein ChrimsonR. The patient, whose visual function was previously limited to light detection, was able to perceive, locate, count, and touch different objects. However, due to the sluggish temporal dynamics of optogenetic proteins, there is likely to be a loss of temporal resolution in optogenetically-transduced RGCs (Fine and Boynton, 2015).

Electronic visual prostheses

Electrical stimulation of neurons Electrical stimulation to create phosphenes, visual percepts that are not caused by the normal activity of the eye in response to light, was initially described by Le Roy (1755). Le Roy, influenced by the report of a physician in Dorchester claiming to have restored sight to a blind man, applied electric shocks to a blind volunteer's head, inducing temporary, vivid and pleasant percepts. Le Roy concluded that the optic nerve was activated in a way similar to its natural activity. Volta (1800) was among the first to investigate the determinants of the perceived flashes resulting from electrical stimulation on his own body. He noticed that the intensity, size, and duration of the phosphenes did not change with the current intensity but depended on the parts of the body where the electrodes were applied.

Injection of electrical current close to a neuron with an extracellular electrode can induce membrane depolarization, and in spiking cells, lead to cell firing (Grill and Mortimer, 1995). The physiological basis of electrical stimulation of neurons is reviewed by Tehovnik (1996); Tehovnik et al. (2006); McIntyre and Grill (2000) and Basser and Roth (2000). Briefly, the electrical activation of neurons corresponds to the depolarization of the cell membrane. Depolarization occurs when charge injection in an electrode induces ions fluxes towards other electrodes. Electrodes can be either faradaic, in which oxidation and reduction occur at the interface between electrode and tissue, or capacitive, with charge accumulation at the electrode-electrolyte interface.

This phenomenon is used in so-called brain-machine interfaces to control neural activity with implanted electrodes, with applications in movement disorders such as Parkinson's disease (Coffey, 2009; Hickey and Stacy, 2016), epilepsy (Bergey, 2013), motor control upon spinal cord injury (Ho et al., 2014), neurogenic bladder (lack of bladder control) (Rijkhoff et al., 1997; Lay and Das, 2012), chronic pain (Plow et al., 2012), psychiatric disorders (Williams and Okun, 2013) and restoration of sensory function such as hearing with cochlear implant (Wilson and Dorman, 2008; Carlson et al., 2012; Shepherd et al., 2013). The potential of these therapies is well represented by the success of cochlear implants, with approximately 736,900 recipients worldwide as of December 2019 according to the US National Institute on Deafness and Other Communication Disorders.

Electronic visual prostheses are devices designed to electrically stimulate neurons of the visual system to recover useful vision (figure 1.5). The elicited phosphenes are colored light forms that span a range of geometrical shapes such as "blobs", "lines", and "wedges", that the patients learn to interpret (Zrenner et al., 2007; Sinclair et al., 2016). Visual prostheses can be classified depending on the stage of the visual pathway they target (for a review of the recent clinical trials for visual prostheses, see Mirochnik and Pezaris (2019)).

Retinal prostheses The retina is an ideal target for the electrical restoration of vision. Indeed, even in advanced stages of RP and AMD, sufficiently many RGCs survive to elicit visual perception. The retina is relatively easy to access compared to locations further down the visual stream, and its laminar structure simplifies the design of the implant. Finally, retina stimulation preserves the high-level information processing occurring downstream in the visual pathway. However, retina implants are limited to cases where the remaining visual system is intact, which excludes conditions such as glaucoma or traumatic injuries. I will describe in



Figure 1.5: Electronic visual prosthetics designs (adapted from (Zrenner, 2013) and Airaghi Leccardi (2020)) A. Retinal prostheses. B. Optic nerve prosthesis. C. LGN prosthesis. D. Cortical prosthesis. E-G. Retinal prosthetics can be implanted either epiretinally (E Fundus photograph of a retina with the Argus II retinal implant Bloch and da Cruz (2019)), subretinally (F The Alpha AMS implant, (Edwards et al., 2017)) or suprachoroidally (G, Ayton et al. (2014)).

more detail the design and function of retinal prostheses in section 1.2.

Optic nerve prostheses The optic nerve was considered as a potential stimulation target due to the ease of access of its extracranial segment, and the fact that it does not require optical transparency, which would make it effective in cases of eye trauma. Moreover, compared to a retinal or cortical implant with the same number of electrodes, an optic nerve prosthesis would allow a wider area of the visual field to be recovered. Implants are either made of non-penetrating cuff electrodes around the nerve (Veraart et al., 1998) or penetrating electrodes (Brelén et al., 2006). Evoked potentials via the optic nerve have been shown to have the same waveshape as normal visual potentials (Brelén et al., 2010; Delbeke, 2011) and to elicit spatially selective activation of the visual cortex (Gaillet et al., 2020). With training, previously blind patients could recognize and orient complex shapes (Brelén et al., 2005; Veraart et al., 2003) and perform object localization, discrimination, and grasping (Sakaguchi et al., 2009; Veraart et al., 2003), although very slowly.

LGN prostheses The lateral geniculate nucleus (LGN) receives inputs from the retina and projects to the primary visual cortex. Given its compact structure and retinotopic organization, it is a potential target for wide-field vision restoration in diseases where the optic nerve is not functional, such as glaucoma, diabetic retinopathy, or traumatic injuries (Pezaris and Reid,

1.1. Retinal degeneration

2007; Kyada et al., 2017). Interestingly, the fovea is overrepresented in the LGN, which could allow for higher resolution vision compared to retinal stimulation (Pezaris and Eskandar, 2009; Panetsos et al., 2011). Moreover, the color channels are anatomically distinct in the LGN (Wiesel and Hubel, 1966), making it easier to selectively activate them, which may allow restoring color vision (Pezaris and Reid, 2007; Panetsos et al., 2011). A study in the macaque showed by analyzing saccades in response to LGN stimulation that phosphenes could be elicited (Pezaris and Reid, 2007). Moreover, it has been shown in rats and rabbits that cortical responses to LGN stimulation were similar to the activity evoked by natural vision (Panetsos et al., 2011).

Cortical prostheses The visual cortex has many advantages as a stimulation target (see Najarpour-Foroushani et al. (2018) for a recent review). First, contrary to retinal prostheses, it can be used for vision restoration in all forms of blindness other than cortical injury or stroke. Second, its large surface area, combined with the cortical magnification factor, makes it possible to implant a large number of electrodes in cortical areas subserving central vision (Daniel and Whitteridge, 1961; Harvey and Dumoulin, 2011), potentially offering a higherresolution visual experience than LGN or retinal implants. Third, the stereotactic implantation of electrode arrays is a relatively straightforward procedure compared to implanting deep LGN electrodes or microarrays onto, or under the retina. Research on visual cortex electrical stimulation started with the observations of Löwenstein and Borchardt (1918), who stimulated the occipital cortex of soldiers with an occipital bullet wound. In 1929, Foerster (1929), quickly followed by Krause and Schum (1931), demonstrated that occipital cortex stimulation could elicit stable, punctate phosphenes (see Donaldson and Brindley (2016) and Lewis and Rosenfeld (2016) for historical reviews). Brindley and Lewin (1968) developed the first visual prosthesis, using electric stimulation of the occipital cortex in blind subjects. Their device was made of 80 $1mm^2$ surface platinum electrodes embedded in a silicon substrate and molded to the recipient's occipital cortex. These electrodes were tethered to a subcutaneously-implanted set of radio receivers, allowing wireless control of individual electrodes. The prosthesis was implanted in a 52-year old woman who had lost her vision as a result of bilateral glaucoma. 32 electrodes could evoke individual visual phosphenes, and simultaneous stimulation of adjacent electrodes was described as evoking phosphenes that could be combined to convey patterns such as letters (Donaldson, 1973).

In 1978, the Dobelle (2000) cortical prosthesis, made of surface electrodes, provided limited object recognition and mobility to its recipients with only 21 phosphenes. Notably, the participant gradually learned to recognize letters (VA of about 20/1200) and count fingers.

Recently, Chen et al. (2020) validated a 1024-channel prosthesis made of intracortical electrodes in the monkey. V1 stimulation was able to convey shape perception, and successive activation of electrodes could elicit motion perception.

Several clinical trials are ongoing for cortical devices. Second Sight (Second Sight Medical Product Inc., Sylmar) has started a six-subjects early feasibility study of its Orion Cortical Visual Prosthesis System (NCT03344848), and the Universidad Miguel Hernandez de Elche in Spain has begun a clinical study on the CORTIVIS device (NCT02983370). Many challenges remain to be addressed to provide implanted patients with useful artificial vision. Current designs rely on head-scanning of the scene and do not make use of eye movement. Similar to the retina, the

cortex may experience reorganization after blindness (see section 1.1.2) that may complicate signal encoding and phosphene mapping.

Cognitive prostheses and sensory substitution

Sensory substitution is the principle of compensating for the loss of a sensory modality by stimulating another functional sense with an artificial sensor. These devices are also denoted as cognitive prostheses (Liu et al., 2018). They can either make use of the tactile sense (with devices as simple as a mobility cane (Bach-y Rita and W. Kercel, 2003), vibro- or electro-tactile stimulation of the skin or tongue (Bach-y rita et al., 1969; Bach-y Rita et al., 1998; Bach-y Rita and Collins, 1971; Deroy and Auvray, 2012)) or auditory sense (Capelle et al., 1998; Hanneton et al., 2010; Meijer, 1992; Liu et al., 2018). Liu et al. (2018), for example, developed an augmented reality device that identifies important aspects of the visual scene such as obstacles or objects, and describe them verbally to the user.

1.2 Vision restoration with retinal prostheses

1.2.1 Prostheses designs

Image capture

Two main methods are currently in use to capture and send the visual inputs to the intraocular component of retinal prostheses. The first one, denoted "intrinsic image capture", relies on the patient's eye to focus the image on a photovoltaic component positioned in the eye. The photovoltaic part, either a microphotodiodes array (MPDA) (Lorach et al., 2015), conjugated polymers or organic semiconductors coupled with titanium nitride cathodes (Ghezzi et al., 2011, 2013; Antognazza et al., 2016; Maya-Vetencourt et al., 2017; Chenais et al., 2021a; Ferlauto et al., 2018), converts incident light rays into electric pulses. This method takes advantage of the natural light-mapping on the retina for impulse generation. In practice, however, natural light focused on photovoltaic interfaces only drives currents in the order of nanoamperes, while the inner retina neurons have activation thresholds of 10mA (Stett et al., 2000; Zrenner, 2002). The alpha-IMS implant circumvented this problem by supplying an external power source, which amplifies the small currents generated by the MPDA sufficiently to activate the inner retina while retaining the retinotopic organization of the stimuli (Degenaar et al., 2009; Zrenner et al., 2011).

The second method, "extrinsic image capture", relies on an external video camera to capture the visual scene. Images are then processed in real-time by a small portable computer (the visual processing unit (VPU) in the Argus II retinal implant) and converted into electrical signals in an electrode array or projected on a photosensitive component with pulses of nearinfrared light. One of the main criticisms of using an external imaging system is that image capture occurs independently of eye position, whereas photosensitive arrays can take advantage of natural eye movements.

Implant positioning

Epiretinal implants are positioned on the inner side of the retina, in contact with the GCL, whereas subretinal implants are set in place of the photoreceptor layer. Epiretinal stimulation activates RGCs with submillisecond pulses, whereas subretinal stimulation activates cells in the INL, as photoreceptors naturally do, with longer (1-10 ms) pulses. The limited available space allows subretinal prostheses to be held in place without a tack, contrary to epiretinal prostheses (Hadjinicolaou et al., 2015; Rizzo, 2011). Additionally, network-mediated stimulation of RGCs is expected to preserve some features of natural image processing within the retina (Lorach et al., 2015). However, subretinal placement of the implant is more challenging (Koitschev et al., 2015) and presents an increased risk of retinal detachment (Kitiratschky et al., 2015).

Finally, the suprachoroidal retinal prosthesis is implanted between the choroid and sclera. The position of the suprachoroidal prosthesis intends to avoid retinal damage that can be caused by direct contact (Hadjinicolaou et al., 2015). However, the distal location of the stimulating contacts to the retina requires higher stimulation currents than with other approaches (Bareket et al., 2017), which increases the risk of damage (Kanda et al., 2004) and the impact of current spreading, hence limiting visual outcomes (Fujikado et al., 2016). Both the epiretinal (Sekirnjak et al., 2006) and suprachoroidal (Kanda et al., 2004) implants have been shown to directly stimulate the retinal ganglion cells predominantly, while the subretinal implants evoke retinal ganglion cells (Stett et al., 2000).

Power and data transmission

Wireless power and data transfer offer significant advantages for a retinal prosthesis. A wireless implant allows for minimally invasive surgery and unrestricted eye movement, preserving the natural oculomotor behavior (Hafed et al., 2016). However, given the large number of electrodes likely to be implanted, it is a significant challenge for a wireless interface to transmit data signals and provide enough power to the stimulating hardware. The first wireless implants used electromagnetic induction, where an alternating electric current in a wire coil (the transmitter) induces a magnetic field which causes current to flow in a separate coil (the receiver). However, the receiver was still wired to the electrode array, hampering eye movement. Even if the wire is attached to the eye, the energy transfer efficiency varies with the relative position of the two coils (Ng et al., 2011). Moreover, absorption of electromagnetic energy by tissues, which increases exponentially with frequency, imposes an additional constraint to prevent tissues from overheating (Al-Kalbani et al., 2012). New generation implants such as POLYRETINA or PRIMA (see figure 1.9), which use photovoltaic components, receive both visual information and energy in projected light. However, this elegant solution exposes to risk of cornea overheating and thus requires eye-tracking sensors in the goggles to ensure real-time adjustment of the beam for it to enter the pupil without irradiating the cornea (Chenais et al., 2021a).



Figure 1.6: Normalized phosphene brightness (*A*) and size (*B*) as a function of amplitude averaged across nine electrodes in an Argus I retinal implant user. Data are fitted using linear regression. *C*. Predicted percepts with increasing amplitude. Adapted from Nanduri et al. (2012). The brightness and size of the phosphenes is approximately proportional to currents amplitude.

1.2.2 Prostheses function

Controlling phosphenes

Phosphenes evoked by a single electrode are generally found to be small, mostly round or punctate, fuzzy-edged, white-to-yellowish in color, and static. They are characterized by a relatively consistent region of the visual field determined by the retinotopic location of the corresponding electrodes. Given these characteristics, phosphenes are usually considered to be the pixels of artificial vision, and individual phosphenes are used as the elementary building blocks to compose a visual scene.

Modulation of stimulus parameters such as pulse amplitude, duration, count, and frequency has been found to affect phosphene characteristics, allowing modest control over size, brightness, duration, or color (e.g. Greenwald et al. (2009); Nanduri et al. (2012); Rizzo et al. (2003a); Horsager et al. (2009); Klauke et al. (2011); Lee et al. (2013)). Nanduri et al. (2012) for example showed that the brightness and size of perceived phosphenes in one Argus I user scaled linearly with pulses amplitude (figure 1.6) and frequency (figure 1.7).

The extent of control differs depending on the device. Devices in which natural light directly activates the implant provide minimal control. Photovoltaic devices such as Alpha AMS only support basic contrast and gain modulation, without control over pulses' shape or polarity, for example (Stingl et al., 2015, 2017).



Figure 1.7: Normalized phosphene brightness (**A**) and size (**B**) as a function of frequency averaged across nine electrodes in an Argus I retinal implant user. Data are fitted using linear regression. **C**. Predicted percepts with increasing frequency. Adapted from Nanduri et al. (2012). The brightness and size of the phosphenes is approximately proportional to currents frequency.

Implant tuning

After undergoing surgery, the stimulation protocol of extrinsic epiretinal image capture implants is calibrated in a lengthy process by determining the perceptual threshold of each electrode individually, that is, the minimal amount of charge required to reliably elicit a phosphene (see for example Hornig et al. (2008) or Finn et al. (2018)). The subject is asked to report the presence of phosphenes for increasing current intensity or frequency levels, with some inserted "placebo" stimulations where no current is injected, but a sound indicates stimulation.

For the IMI retinal implant, made of 49 electrodes, the procedure took at most 45 minutes by only tuning currents amplitudes. Moreover, alterations in phosphene thresholds over time (Davis et al., 2012) may require this process to be repeated to ensure a consistent visual experience. Given the exponentially increasing number of electrodes, as well as the large number of parameters that can be tuned (pulse polarity, shape, amplitude, and frequency, for example), there is a need for faster, automated tuning procedures.

Image translation

The use of an external image-capturing system allows using image processing algorithms to adapt the electrode array activity or the light pattern. The classical approach to shape the stimulation pattern is simply to pixelate a static image, which is then linearly converted to an electrical signal with a fixed pulse shape, but varying amplitude or frequency. According to Dobelle (2000), a critical factor in gradually restoring visual function in a blind individual with a cortical prosthesis (see section 1.1.3) was the use of image processing algorithms providing edge detection and magnification. Barnes et al. (2016), for example, showed that the use of a vision processing algorithm called Lanczos2 filtering to avoid artifacts due to downsampling led to better visual performance in suprachoroidal prosthesis users. With the Argus II device, enhancing the outlines of the target improved patients' performance in shape and object recognition tasks (Luo and da Cruz, 2016). Automatically extracting relevant shapes from the visual scene while removing the background and present it to the patient in a zoomed-out view also appeared to improve performance in a face detection task (Stanga et al., 2013). Barnes et al. (2016) demonstrated that a Nyquist bandlimited downsampling filter improved the performance of patients equipped with a 20 electrodes suprachoroidal prosthesis in a light localization task, compared to minimal vision processing or regional averaging filtering. McCarthy et al. (2011) proposed to use an image segmentation algorithm that detects ground surface and optimizes the corresponding pixels to help identify ground surface and help obstacle avoidance. Recently, Han et al. (2021) assessed a segmentation algorithm as a preprocessing strategy to simplify the visual scene in a simulation of prosthetic vision with Argus II, showing that object segmentation may better support scene understanding than image processing based on visual saliency and monocular depth estimation.

Another image processing strategy is to provide the patient with visual semantic inputs using symbolic cues, making use of the limited number and quality of phosphenes to convey maximal utility percepts. For example, encoding of letters using "visual Braille" allowed a cortical prosthesis user (Dobelle et al., 1976) and Argus II users (Lauritzen et al., 2014) to read. More generally, a camera could automatically recognize environment features such as faces or obstacles and encode them in visual cues. This strategy was evaluated by Parikh et al. (2013), who used saliency-based cueing to help participants in mobility and search tasks using simulated prosthetic vision. A blinking phosphene attracted participants' attention and improved their performance in object location, object avoidance, and reduction in head scanning. Similarly, Mohammadi et al. (2012) proposed the use of a range-finding algorithm to estimate the distance to objects, which would be indicated to the user with a group of phosphenes reserved for this purpose. The IRIS-IMI implant was initially designed to include an image processing algorithm mimicking the center-surround organization and temporal resolution of midget and parasol RGCs (Eckmiller et al., 1999, 2005; Hornig et al., 2008). The tuning procedure of the filters' parameters was based on feedback from users, asked to select the parameters maximizing perceptual resolution in a sequential optimization. However, there was no report concerning the impact of this image processing on visual performance.

Rehabilitation and perceptual learning

Patients fitted with a prosthetic device have to engage in a training and rehabilitation program in order to learn how to deal with poor resolution and contrast, restricted field of view, high-eccentricity, geometric distortion, and other limitations. Learning supervised by clinicians, family and friends play a major role in visual rehabilitation (Chen et al. (2009b) discuss in great detail rehabilitation programs for implant users). Patients have to learn to interpret the phosphenes, to use head-scanning, to develop visuomotor coordination and eccentric read-



Figure 1.8: An example of image processing in the Argus II retinal implant (Bloch and da Cruz, 2019). The image first undergoes edge detection and is converted to greyscale. It is then transformed by circular binarisation into a black-and-white pixelated format. Finally, the image is inverted, with the addition of four greyscale levels.

ing. Once the patient has learned to use this new artificial vision, visual performance can be assessed. Notably, most clinical trial participants had bare light perception or less prior implantation, with a few including those who are able to identify hand movements.

The behavioral manifestation of cortical plasticity is perceptual learning, i.e. long-lasting improvement in behavioral performance in a perceptual task (Sagi, 2011). The psychophysical literature on visual perceptual learning (reviewed in Beyeler et al. (2017b)) shows that improvement on tasks performed on basic stimulus attributes such as grating detection, orientation tuning, contrast discrimination, or visual acuity is extremely slow and does not generalize to other tasks. On the contrary, more considerable perceptual learning occurs for more complex tasks.

Perceptual learning in prosthetic users may be very different compared to perceptual learning in sighted subjects, given the very different nature of prosthetic vision as experienced by patients (Erickson-Davis and Korzybska, 2020). However, consistent with findings on sighted subjects, patients implanted with retinal prostheses generally show barely any improvement on simple perceptual tasks, such as contrast sensitivity (Castaldi et al., 2016) or motion discrimination (Dorn et al., 2013; Castaldi et al., 2016), even after extensive training over several months. Only detection thresholds slightly decreased as a function of time since surgery in an Argus II user (Castaldi et al., 2016). The most significant perceptual improvements found in prosthetic users have generally been reported in more complex tasks such as moving in space, locating a sizeable bright square on a screen, and identifying large-print letters (Chader et al., 2009; Zrenner et al., 2011; Humayun et al., 2012; Da Cruz et al., 2013; Dorn et al., 2013; Stingl et al., 2013; Rizzo et al., 2014). However, there is little evidence that this is due to distortions becoming less perceptually apparent. Instead, most improvements seem to be relatively taskspecific and are the result of patients becoming better at interpreting distorted input. This suggests that the transfer of performance improvement outside controlled clinical settings may be limited. As a consequence, in order for patients to recover useful visual function in their everyday life, rehabilitation programs should incorporate diverse tasks and contextual environments (Beyeler et al., 2017b).

Visual performance evaluation

A key concern in the field of artificial vision is the lack of standardized procedure to quantify implants performance, limiting the ability to compare results (Rizzo et al., 2014; Ayton
et al., 2020). As discussed in detail by Dagnelie (2008), prosthetic vision should be assessed in a variety of contexts and tasks: from simple light and motion perception to the ability to perform activities of daily living. Recently, the Harmonization of Outcomes and Vision Endpoints in Vision Restoration (HOVER) Task Force addressed this issue by providing detailed methodological recommendations to assess visual function (Ayton et al., 2020). Importantly, prosthetic vision, as reported by users, is qualitatively and functionally different compared to natural vision. The fact that current implants are more akin to sensory substitution instead of replacement or restoration and that the use that patients make of artificial vision is highly variable, even with the same device (Erickson-Davis and Korzybska, 2020) should be carefully taken into account when designing evaluation tasks. For example, some Alpha-IMS users who had good performance in the standardized test were not able to use their implant in daily life, whereas others with worse performance in these tests could (Stingl et al., 2013).

A crucial aspect that is usually not measured is the ease of use of prosthetic vision. Indeed, in a study on Argus II and IRIS II, the use of the device was reported as being cognitively very fatiguing, even years after implantation (Erickson-Davis and Korzybska, 2020). Even if an implant provides some rudimentary vision, patients could stop using it because of its cognitive requirements.

Visual outcomes for patients

I will now cover some results concerning the most noticeable implant development projects. A list of present and past projects is available at http://www.eye-tuebingen.de/zrenner/ retimplantlist/. Many of them were discontinued due to the lack of significant enough benefits for the patients.

Visual outcomes with Argus II and IRIS II The Argus II (Second Sight Medical Product, Inc., Sylmar. Argus Panoptes ("all-seeing") is a many-eyed giant in Greek mythology) is an epiretinal prosthesis featuring a 60-electrode array covering a 20° field of vision (Da Cruz et al., 2013; Humayun et al., 2012; Dorn et al., 2013). It received Conformité Européenne (CE) marking in March 2011 and the Food and Drug Administration (FDA) approval in February 2013. The external parts of the device include a glasses-mounted video camera, a portable computer for processing the images, and an external coil, built into the sidearm of the glasses, for wireless communication using radio-frequency (RF) telemetry and induction of power. The internal components consist of an internal coil that acts as a receiver, an application-specific integrated circuit (ASIC) that converts the RF into electric signals stimulation, and the electrode array, held in place by a retinal tack. The internal coil and ASIC are sealed in a protective hermetic casing, which is placed on the surface of the globe, connected with a transscleral cable to the intraocular array. The implantation surgery takes about two to three hours. Among the thirty patients enrolled in the five-year clinical trial (Humayun et al., 2012; da Cruz et al., 2016), conjunctival erosion is the commonest complication (10%). During the trial, one patient's device had to be explanted due to recurrent erosions, and one patient suffered from retinal detachment, which was successfully repaired. Two patients developed endophthalmitis, but all were treated successfully with intravitreal antibiotics, and they retained good functional use of their device. Overall, 60% of patients had no serious adverse events, and the device appears to be

well tolerated and safe.

All patients reported reliable perception of phosphenes. The device enabled more than half of the recipients (15/28) to identify the direction of motion of a high-contrast bar moving in varying directions on a flat LCD screen (Dorn et al., 2013). Some of the patients were able to localize squares (Ahuja et al., 2011) or even discriminate different geometric shapes (Da Cruz et al., 2013) when presented with high-contrast targets on a flat LCD screen. After practice, single-letter recognition seems to take somewhere between a few seconds to three and a half minutes due to prolonged head scanning, and four subjects were able to consistently read unrehearsed short words of up to four letters (Da Cruz et al., 2013). The best grating visual acuity, expressed in the log of the minimal angle of resolution in minutes of arc (logMAR), was logMAR 1.8 (Snellen equivalent of 20/1262) from worse than logMAR 2.9 pre-operatively (Humayun et al., 2012). Orientation and mobility functions were tested by following a white line on a dark floor and locating a dark door on a white wall from the center of a room. Patients performed significantly better with the device switched on (Humayun et al., 2012).

IMI GmbH, which developed the epiretinal implant IRIS (Intelligent Retinal Implant System) with promising results (Hornig et al., 2008), has been acquired by Pixium Vision in 2007. Pixium then developed the IRIS 2 device, which uses a 150-microelectrode array and has received a CE Mark in 2016 (Gabel, 2017). It uses a neuromorphic event-based camera, which mimics the temporal resolution of the retina (Lorach et al., 2012), and like IRIS, uses an image processing algorithm mimicking retinal processing that is tuned based on patients feedback (Eckmiller et al., 2005). In 2017, the company announced promising results regarding visual function for the 10 subjects enrolled in a clinical trial. However, the clinical trial ended in 2019, apparently due to concerns regarding the lifespan of the device.

Visual outcomes with Alpha IMS and Alpha AMS The alpha-IMS is a 3 mm× 3 mm subretinal implant made of an MPDA connected with a cable to an internal coil for supplementary inductive power supply (Zrenner et al., 2011). It does not rely on an external imaging system. The array is made of 1500 photodiodes, each connected to a 50 μ m x 50 μ m titanium nitride microelectrode via an amplifier. The implant covers a visual field of 11° × 11°. Implantation requires a 6 to 7 h surgery. The Alpha IMS device was first trialed on nine patients with retinitis pigmentosa and cone-rod dystrophy and received a CE Mark in July 2013. One was excluded due to complications during the implantation procedure. 7/8 recipients were able to localize a light source, 5/8 were able to detect motion, and 6/8 were able to perform grating acuity testing (Stingl et al., 2013). Most importantly, 5/8 recipients reported functional use of their implant in their daily life, which they reported as being the most important and rewarding aspect. Two participants could not make use of their artificial vision in daily life due to perceptual fading.

The Alpha-AMS device (Retina Implant AG, Germany) is the successor of Alpha-IMS, with a 1600-sites MPDA. The size of the electrode area of the Alpha AMS chip is 2.8 × 2.8 mm², corresponding to a visual field of 9.3°. Via a handheld device, the patient can adjust the brightness and contrast of the signal to the surrounding brightness using two controllers. This device displayed at least as well performances compared to Alpha-IMS. The highest visual acuity measured was Snellen 20/546 (logMAR 1.43), an important improvement compared to previous devices (Stingl et al., 2017; Edwards et al., 2017).

Visual outcomes with PRIMA The only subretinal prosthesis currently undergoing clinical testing is the PRIMA device by Pixium Vision (NCT03333954, NCT03392324) (Lorach et al., 2015; Wang et al., 2012). It receives operating power and digitally processed visual data from a light projector carried on video goggles (powered from a belt pack). Near-infra-red pulsed light (880 nm) is used to avoid the perception of light by the remaining photoreceptors. The implant comprises 378 microphotodiodes over a 2×2 mm, 30 µm thick array, of which multiple devices can be injected (Lorach et al., 2015; Palanker et al., 2020). Each microphotodiode is 100µm in diameter. A clinical study was performed on five patients (NCT03333954) with geographic atrophy, showing that central VA could be restored up to 20/460 (1,37 logMAR) (Palanker et al., 2020), only 0.07 logMAR worse than the sampling limit for 100µm pixels (1.3 logMAR). The two other patients for which implant positioning was successful recovered central VA of 20/550 (1.44 logMAR) and 20/500 (1.4 logMAR). Patients also reported flicker fusion at frequencies exceeding 30 Hz. In patients with AMD, peripheral vision is preserved, so the investigators intend to use PRIMA with transparent glasses acting like augmented reality goggles to simultaneously use the prosthetic central and natural peripheral vision.

In all patients, the prosthesis was implanted successfully under the macula, although in two patients, it was implanted in unintended locations: within the choroid and off-center by 2mm. All five patients could perceive white-yellow visual patterns with adjustable brightness in the previous scotomata. The three participants with optimal placement of the implant demonstrated prosthetic acuity of 20/460 to 20/550.

Visual outcomes with suprachoroidal implants Currently, no suprachoroidal prostheses have been approved beyond phase I clinical trials (Fujikado et al., 2016; Ayton et al., 2014). Ayton et al. (2014) successfully implanted a device that allowed late-stage RP patients to localize light and recognize letters, although with an optotype acuity of approximately 20/8000.

1.2.3 Limiting factors for vision restoration

The theoretical visual acuity (minimal angle of resolution θ) that would be achieved if performance was limited by the spacing of electrodes is : $\tan \theta = \frac{\text{pixel pitch}}{\text{focal length}}$ (Daschner et al., 2018). In practice, however, only a handful of patients show performance close to this theoretical expectation (see section 1.2.2). For example, the theoretical MAR with the Alpha AMS chip is logMAR = 1.15, which corresponds to a Snellen fraction of 20/280 (Daschner et al., 2018), but in practice, the two best patients were able to distinguish Landolt C-rings of 20/1111 and 20/546 (Stingl et al., 2017).

Insights from electrophysiological and psychophysical experiments as well as theoretical considerations (reviewed in Fine and Boynton (2015)) suggest that this poor performance may be due to interactions between implant electronics and the underlying neurophysiology of the retina. Recent progress in the understanding of electrodes/tissue interaction and information processing in the retina suggests that these issues can be overcome and that functional vision restoration is not pure quixotism. I will now cover the main challenges and the solutions that have been proposed to address them.

Surgery Implantation of the device requires an invasive surgery which can result in adverse effects. The most commonly reported events were elevated intraocular pressure (IOP), conjunctival erosion, and retinal detachment (Edwards et al., 2017; Kitiratschky et al., 2015). However, recent wireless implants seem to require much simpler surgeries compared to previous devices (figure 1.9).

Biocompatibility, charge injection limit and hermiticity One of the key obstacles to developing a visual prosthesis is the observed deterioration of the interface between electrodes and neural tissue (see Polikov et al. (2005) for a comprehensive review). Indeed, chronic electrode implantation for both recording and stimulation revealed the formation of a glial scar due to mechanical trauma of insertion and a chronic inflammatory response corresponding to foreign body reaction (see, e.g., McCreery et al. (2010)). Moreover, the extracellular fluid contains a high concentration of ions which can harm electrical circuits, and adequate protection must be integrated to provide hermiticity for implanted components.

Irreversible electrochemical (faradaic) reactions occurring at the electrode/tissue interface may lead in particular to electrode dissolution due to the oxidative formation of soluble metal complexes, in addition to water electrolysis causing gas bubble formation and injurious pH shifts within surrounding tissue. For this reason, typical pulses used in neural stimulation have a biphasic waveform, where charges injected during the cathodic phase and the anodic phase compensate (charge-balance), resulting in a zero-net charge injection. Charge-balance limits irreversible reduction and oxidation reactions (Cogan, 2008), but they can still occur with high charge density (McCreery et al., 1990, 2010). Charge density itself depends on the amount of electric charge delivered per stimulus pulse (charge per phase) and the electrode surface area. The risk of irreversible electrochemical reactions can thus be lowered by using electrodes with higher surface areas (McCreery et al., 2010; Negi et al., 2012), although this may be at the expense of reducing the stimulation resolution. Another possibility is to use materials with high charge injection capacity, such as iridium oxide films (Negi et al., 2010), poly(3,4ethylenedioxythiophene) coating (Wilks et al., 2009), roughened silicon coated with platinum (Negi et al., 2012) or silicon electrodes containing embedded carbon nanotubes (Musa et al., 2012).

In general, an excess charge delivery can damage the neuronal tissue, even in the absence of faradaic reactions (Bartlett et al., 1977; Brummer et al., 1983). Indeed, Bullara et al. (1988) observed that neuronal loss was independent of electrode type (i.e., faradaic vs. capacitative), suggesting that the phenomenon can occur in the absence of electrochemical reactions occurring at the electrode/tissue interface. The authors hypothesized that the damage might be mediated by stimulation-induced neuronal hyperactivity. Support for this idea was provided by administering an NMDA receptor antagonist during stimulation of the cat cortex, which reduced the degree of neuronal damage compared to untreated animals and suggested a glutamate-mediated excitotoxicity (Agnew et al., 1993).

Together, these phenomenons can lead to gradual electrode failure over months to years (Hochberg et al., 2012; Rousche and Normann, 1998) or an increase in stimulation currents thresholds to excessive levels (Davis et al., 2012). Davis et al. (2012), for example, investigated the ability of an electrode array to elicit phosphene two years after implantation into the visual

cortex of a macaque. 77/96 individual electrodes failed to consistently produce behavioral responses at currents up to 200 mA. Making use of biocompatible materials is thus a major issue (Montezuma et al., 2006; Zhao et al., 2018), as well as designing devices that are robust to electrodes malfunction, for example, by allowing reconfiguration of stimulus parameters to maintain device efficacy over time.

Visual field The normal monocular human visual field extends from approximately 60° nasally to 107° degrees temporally from the vertical meridian, and roughly 70° degrees above and 80° below the horizontal meridian. The binocular visual field is the superimposition of the two monocular fields. It extends on about 180° in the horizontal meridian. The macula corresponds to the central 17° diameter of the visual field; the fovea to the central 5.2° , and the foveola to $1^{\circ} - 1.2^{\circ}$ diameter. The potential breadth of the induced visual field is directly proportional to the two-dimensional area of the electrode array that stimulates the retina. Current commercially available implants cover less than 20° in the peripheric visual field. The use of larger arrays will require larger and potentially riskier incisions to insert the device. Several strategies have been proposed to address this challenge. The PRIMA device, for example, can be inserted successively in several components. A wide-field epiretinal photovoltaic retina implant was recently developed (Ferlauto et al., 2018). Made of 2215 stimulating photodiodes, the POLYRETINA implant (see figure 1.9) is thin and foldable, which makes it implantable through a small scleral incision, and covers a visual field of 46.3 degrees.

Perceptual fading Another concern is the tendency for perceptual fading; that is the lack of image persistence beyond a few seconds, for both camera-based (Pérez Fornos et al., 2012) and photodiode-based prosthetics users (Hafed et al., 2016; Stingl et al., 2015; Zrenner et al., 2011). As a consequence, for users to continuingly see an object, they need to "refresh" the image by shaking their head. This is reminiscent of the progressive disappearance of percepts when the retinal image is artificially stabilized (Tulunay-Keesey, 1982), and perceptual fading in prosthetic users has thus been postulated to reflect a lack of naturally occurring microsaccades (Weiland et al., 2016; Hafed et al., 2016). Indeed, human eyes undergo small involuntary movements called microsaccades when fixating on an object. Among other functions, microsaccades are involved in the refreshing of retinal images, preventing them from fading (Coppola and Purves, 1996; Riggs et al., 1953; Martinez-Conde et al., 2006; Intoy and Rucci, 2020). However, RP patients using the subretinal implant Alpha IMS seemed to experience perceptual fading (Hafed et al., 2016; Stingl et al., 2015) even though they exhibited normal saccadic and microsaccadic behavior. AMD patients using the PRIMA subretinal devices did not perceive fading (Palanker et al., 2020). However, it is interesting to note that two of the Argus II patients perceive constant images without having to consciously shake their heads (Pérez Fornos et al., 2012).

It is not clear why visual fading occurs in some subretinal implant users despite normal oculomotor behavior. Fading may also depend on the extent of electrodes' projective fields (the set of neurons each electrode influences). If the extent of microsaccade-induced stimulus motion on the retina is smaller than electrodes projective fields, the same neurons will stay

1.2. Vision restoration with retinal prostheses



Figure 1.9: A-E. The PRIMA retinal implant (adapted from (Lorach et al., 2015) and the Pixium Vision press media kit). A-B. Each pixel consists of two to three (shown here) photodiodes connected in series between the central active (1) and surrounding return electrode (2). Scale bar: 50µm. C. Picture of a PRIMA consisting of 205 electrodes at a pitch of 140 µm. D. Fundus image from one non-human primate bearing the PRIMA two weeks after surgery. E. Location of the subretinal photovoltaic implant, as seen by optical coherence tomography of the retina. F-I. The POLYRETINA implant (adapted from Ferlauto et al. (2018)). F. Scanning electron microscopy image (40° tilted view) of a photovoltaic pixel of the POLYRETINA implant. Scale bar is 10µm. G. Picture of POLYRETINA. Four anchoring wings with holes are present for attaching the prosthesis with retinal tacks. H. POLYRETINA folded before injection. I. 3D model after epiretinal placement.

activated in response to a constant visual stimulus (Weiland et al., 2016).

This may explain why PRIMA users did not experience fading. One of the underpinning mechanisms of perceptual fading is the desensitization of stimulated cells following repeated electrical stimulation, and the time dynamics of the two closely match (Freeman et al., 2011). When directly stimulated, RGCs can follow stimulation pulse rates frequencies over 200 Hz (Fried et al., 2006; Sekirnjak et al., 2006). However, the response of RGCs to network-mediated stimulation rapidly decreases following repeated stimulation, and this effect is more substantial as frequency increases (Jensen and Rizzo, 2007; Freeman et al., 2011). The mechanism of RGCs network-mediated response reduction may be the intrinsic desensitization of bipolar cells (Walston et al., 2018) combined with inhibition from amacrine cells (Freeman et al., 2011; Walston et al., 2018).

Stimulation strategies have been designed to limit this desensitization. Chenais et al. (2021b) for example, showed with the POLYRETINA device that protocols including spatial and temporal modulation could maintain responses for up to 4.2 s. Spatial modulation corresponded to the stimulation of RGCs with sequential activation of photovoltaic pixels corresponding to different subsections of its receptive field. Temporal modulation included pulse trains with irregular pulse widths and randomized time-varying stationary interrupted pulse sequences.

Head and eye movements A particular challenge for systems dependent on a head-mounted camera is the dissociation between the axis of the camera and the direction of gaze. Normal localization of an object depends on the retinotopic location of the target relative to the orientation of the subject's eye and head. As a consequence, a shift in the eye position creates a proportional shift in the apparent location of an object induced by visual prosthetic stimulation (Brindley and Lewin, 1968; Dobelle and Mladejovsky, 1974; Sabbah et al., 2014). However, eye and head tracking can be implemented in the goggles to adjust image capture depending on eye and head position, which reduces localization errors (Titchener et al., 2018). An intrinsic image capturing implant addresses this issue, as would an intraocular camera (Nasiatka et al., 2005; Hauer et al., 2007).

Nonlinear mapping from RGCs eccentricity to visual field coordinates An additional issue when artificially stimulating RGCs comes from the fact that the mapping of retinal eccentricity to visual field coordinates is nonlinear. Indeed, in the foveola, which contains only photoreceptors, ganglion cell bodies are displaced centrifugally from their cone inputs by several degrees, an effect that extends out as far as 17° (Curcio and Allen, 1990; Watson, 2014). This suggests that prostheses aiming at providing high acuity vision should include an image processing component compensating for this nonlinear mapping.

Electrodes size and current spread Due to their diameter, each electrode activates many cells. Efforts are being made to create devices that can stimulate cells more specifically by using electrodes as small as 9µm in diameter (Sekirnjak et al., 2008) (against 200µm in Argus II). However, as the electrode diameter becomes smaller, the charge density increases for the same supra-threshold stimulating current (Tehovnik, 1996), thereby increasing the risk of tissue damage (see paragraph 1.2.3).

Current spread within the tissue is also an issue, as it leads to cell activation away from the electrode, affecting both spatial and temporal resolution. Indeed, current spread results, for example, in a single pulse applied epiretinally to evoke two action potentials in RGCs. The first spike occurs within 1ms following the stimulation and is due to the direct activation of the cell, and a second spike follows a few milliseconds later, most likely because of the inner retina network activation (Sekirnjak et al., 2006; Fried et al., 2006).

Because of current spread, current densities at a given location are a function of currents injected by all electrodes at any given instant (figure 1.10). This phenomenon, denoted as electrodes cross-talk, is usually considered to decrease the number of effective channels in the electrodes array, strongly affecting contrast sensitivity and spatial resolution (Palanker et al., 2005; Loudin et al., 2007; Wilke et al., 2011; Goetz et al., 2015; Beyeler et al., 2019b). Various strategies have been proposed to reduce the effects of the current spread.

• Return electrodes for focal stimulation: In implants in clinical use today (Argus II and the Alpha AMS), current flows from the activated electrodes to a common return electrode. Return electrodes can be used to confine the spread of current and thus the region of activation. Various return electrode configurations have been investigated in order to maximize the implant's ability to deliver high contrast and high-resolution stimuli (Abramian et al., 2011; Habib et al., 2013; Matteucci et al., 2013; Flores et al., 2016; Fan et al., 2019). The use of local return electrodes



Figure 1.10: Linearity of electrodes crosstalk for simultaneous current injection through pairs of electrodes (from (Jepson et al., 2014a)) **A.** Response curves for an OFF parasol cell (same cell as in B and C) resulting from cathodal current pulses injected through the primary electrode alone (black), from equal amplitude cathodal pulses passed simultaneously through one secondary electrode (teal) and from equal amplitude anodal pulses passed through the same secondary electrode (orange). Vertical lines show the activation threshold for each case. **B.** Layout of primary (white dot) and secondary (black dots) electrodes relative to the electrical image of the targeted RGC (shaded contours). **C.** Measured activation thresholds corresponding to different current ratios for the same electrode pair as in A (inset, black). The black line depicts the linear model fit to the thresholds. Negative secondary current values signify anodal current pulses.

surrounding the stimulating component and isolating it from neighboring "pixels" was shown to greatly improve the restored visual acuity (Lorach et al., 2015) and was later applied in the PRIMA implant (see paragraph 1.2.2).

• Current steering and focusing: Current steering (Mueller and Grill, 2013; Jepson et al., 2014a) or field-shaping (Rattay and Resatz, 2004; Mueller and Grill, 2013) could be used to enhance focalization and elicit more phosphenes than the number of electrodes. Individual ganglion cells can be targeted with the epiretinal electrodes by careful shaping of the electric field based on a computational model of RGCs responses to current (Jepson et al., 2014b). Indeed, simple linear-nonlinear models are able to predict responses of RGCs to arbitrary patterns of multi-electrode stimulation, making it easy to optimize stimulation protocols *in vitro* (Jepson et al., 2014b; Maturana et al., 2016). However, to reach such single-cell resolution, recordings of the electrically-evoked activity are necessary to fit the model. Traditional stimulation strategies consider phosphenes as the pixels of artificial vision. Another approach would be to use interactions between electrodes to shape the RGCs activity pattern globally. This idea was investigated by Spencer et al. (2018), who proposed to use a linear-nonlinear model of RGCs response to electrical stimulation for global shaping of the array activity so as to enhance activation accuracy.

Activation of passing axons Because epiretinal prostheses are in direct contact with the GCL, these devices may accidentally stimulate passing axons from distant ganglion cells (figure 1.11), as suggested by both electrophysiological (Jensen et al., 2003; Behrend et al., 2009; Weitz et al., 2015; Grosberg et al., 2017) and modeling studies (Greenberg et al., 1999; Tsai et al., 2012; Beyeler et al., 2019a). This was interpreted to be the cause of large, arc-shape phosphenes (Fine and Boynton, 2015; Beyeler et al., 2017b) that may severely degrade visual acuity (Rizzo et al.,

2003a,b; Beyeler et al., 2017a, 2019a). Subretinal and suprachoroidal approaches are believed to diminish the risk of passing axons activation due to the increased distance from stimulating electrodes. However, some subjects have reported arcuate and linear percepts with high stimulation levels with both subretinal (Wilke et al., 2011) and suprachoroidal (Sinclair et al., 2016) implants, suggesting that RGCs fiber activation still occurs. Modification of stimulus parameters, especially pulse duration and frequency (Greenberg et al., 1999; Jensen et al., 2003; Fried et al., 2006, 2011; Cai et al., 2013; Weitz et al., 2015; Grosberg et al., 2017) as well as pulse shape have been shown to reduce axon activation. The mechanism of biased activation towards RGCs somas could be preferential bipolar cells activation, which in turn activate RGCs at their soma in a confined region close to the electrode (Freeman et al., 2010; Boinagrov et al., 2014; Weitz et al., 2015). However, this approach eliminates the ability to produce precisely-controlled spike trains in multiple RGC types.

In a computational study, Beyeler (2018) found that appropriate stimulation parameters could bias RGCs activation to the location of the sodium channel band, closed to the soma. Beyeler (2018) simulated a population response to epiretinal stimulation of 0.1-50ms pulse duration in a morphologically and biophysically detailed retinal ganglion cell model. He showed that the excitation thresholds of RGCs were minimal when the electrode was positioned above the sodium channel band, with thresholds rising across the thin segment and eventually plateauing in the distal axon. This suggests that it could be possible to activate neurons without activating distal axons. The "dynamic range", defined as the range of currents activating RGCs somas without activating passing axons, increased with the vertical offset of the electrode. Moreover, as pulse duration increased, the spatial extent of ganglion cell activation in the population was noticeably reduced.

Grosberg et al. (2017) and Tandon et al. (2021) proposed a method to automatically detect passing axons activation in epiretinal prostheses in order to restrict stimulation to cells for which activation is confined to the proximal axon. Finally, Beyeler et al. (2019a) showed using a prosthetic vision simulator that optimal positioning of the implant during surgery could reduce the extent of axon activation.

Using the neural code As mentioned in subsection 1.1.2, despite the significant reduction from 30,000 auditory axons to no more than 22 stimulating electrodes (Sagi and Svirsky, 2018), cochlear prosthetics still provide speech recognition for the severely deaf (Shannon, 2012). This success has largely been driven by advances in signal preprocessing and encoding, particularly leveraging the ability to mimic anatomic frequency encoding in the cochlea by spectral encoding across electrodes rather than by significant advances in device architecture. In the field of visual prosthetics, signal encoding strategies are still at an embryonic stage. Computational modeling suggests that simultaneous activation of the ON and OFF-pathways leads to severe information loss and percept distortion compared to specific only-ON or only-OFF activation (Golden et al., 2019) (see figure 1.12). Interestingly, increased electrode density and reduced current spread only bring marginal improvements when stimulation is unspecific, suggesting that mere improvement in the hardware will not bring additional benefits. Moreover, optimal learning of a linear decoder was shown to substantially reduce the decoded percept distortion (Golden et al., 2019). As discussed in section 1.1.2, whether the brain will be able to make sense



Figure 1.11: A-D Bidirectional propagation of electrically evoked responses due to passing axons activation. Adapted from (Grosberg et al., 2017). A. Fluorescence image showing the density and arrangement of RGC axon bundles with respect to the electrode array. All of the axons run in the same direction toward the optic disc. Arrow indicates the stimulating electrode in panels C-D.B. Electrical images (EIs, spike waveforms of RGCs measured at different electrodes in the absence of electrical stimulation artifact) from three RGCs in a single retina obtained with visual stimulation (no electrical stimulation). Waveforms on the electrodes indicated with numbers are associated with the shaded RGC. C-D. Unidirectional (C) and and bi-directional (D) signal propagation after electrical stimulation at the electrode shown by a red dot. Note similarity of unidirectional (C) image to the EI from the shaded cell shown in B, indicating that this specific cell was activated in isolation. The amplitude of the waveform shown for electrode 6 in all panels was reduced by a factor of 2 relative to the scale bar. E. Trajectory of axon bundles in the human retina, measured from fundus photographs Jansonius et al. (2009). F. Drawings by an Argus II user when asked to depict the phosphenes elicited by individual electrodes activation, represented on top of the corresponding electrode (Beyeler et al., 2019a). G Simulations of perceptual distortions as a result of axonal stimulation, the original image being a Snellen chart(Fine and Boynton, 2015).

of these new inputs is unknown. This suggests that a first approach to enhance perceptual quality is selective, or at least substantially biased stimulation of either ON or OFF-pathways. Modulation of current amplitudes and frequencies (Cai et al., 2011; Guo et al., 2018), as well as pulse polarity (Jensen and Rizzo, 2006) and pulse duration (Im et al., 2018) have also been shown to bias stimulation towards OFF- and ON- pathways.

However, the behavioral consequences of the indiscriminate activation of ON and OFF cells are unknown, even though in some diseases, such as Duchenne and Becker muscular dystrophy, unbalanced retinal signaling starting at early childhood does not necessarily degrade visual performance (Fitzgerald et al., 1994; Zeitz et al., 2015).

Mimicking retinal signal processing would require a much better understanding of the retina. Indeed, among the close to 30 RGCs types identified in the mammalian retina which encode distinct visual features in parallel for transmission to the brain (see 1.1.1), only a subset has been functionally characterized (Sanes and Masland, 2015). Moreover, even if the retina



Figure 1.12: Signal distortion due to unspecific cell types stimulation in simulated prosthetic vision (adapted from Golden et al. (2019)). A. Visual stimulus B-C. Linear reconstruction of the stimulus, using the reconstruction filters obtained from a healthy retina simulation, separately using ON (B) and OFF (C) cells. Note the reverse polarity of the reconstruction for OFF cells. D. reconstruction with all four RGC types combined normally, causing the cancellation of ON and OFF signals. E. Reconstruction with all four RGC types combined, using reconstruction filters obtained with prosthesis stimulation, mimicking optimal learning by the patient.

neural code was perfectly known, identifying the various cell types and specifically stimulate them so as to replicate this code would be a formidable challenge. First, high spatiotemporal precision would be required. In principle, single-cell, single-spike resolution when stimulating RGCs is possible, as it has been demonstrated with high-density epiretinal electrode arrays in isolated peripheral retina in monkey (Sekirnjak et al., 2008; Jepson et al., 2013) and human (Madugula et al., 2020). However, to emulate the neural code, cell types identification and specific stimulation would be required. Cell type identification in an implanted retina is challenging as it usually relies on the measurement of the cells' light-evoked responses. However, Richard et al. (2015) showed that it was possible to use electrical recordings to accurately classify the major RGCs types. Therefore, a "bidirectional" prosthetic device, able to record and stimulate at the same time, could be used for fine-tuned electrical stimulation. The computational challenges in encoding visual stimuli in a bidirectional device are discussed in Shah and Chichilnisky (2020). I will discuss in more detail the strategies that have been proposed for tuning of retina implants in 1.3.1.

Given the differential response of different cell types to electrical stimulation, it could also be possible to specifically stimulate them without having to record the retinal activity (Twyford et al., 2014; Madugula et al., 2020). However, it remains unclear whether the differential sensitivity to electrical stimuli between RGC types is preserved in the degenerated retina.

Instead of directly activating RGCs, one could instead activate the inner retina network to achieve more natural activity. Indeed, studies of the PRIMA device showed that RGCs presented an antagonistic center-surround photovoltaic RF with a size similar to natural RFs, with some cells exhibiting OFF responses (Ho et al., 2018). Moreover, responses to gratings smaller than the cells' receptive field diameter suggested nonlinear summation of subunits (Lorach et al., 2015). These results indicate that this system is able to restore essential features of natural retinal signal processing.

1.2.4 Simulating prosthetic vision

Simulation of prosthetic vision is an invaluable tool to give researchers, regulators, patients, and their families a fair understanding of the outcomes of implantation and to balance benefits and risks. Information provided to the public is often an oversimplified, over-optimistic view of the actual benefits that can be expected from prosthetic vision (as can be seen by visiting the websites of the companies producing these implants). Moreover, prosthetic vision simulations can help designing rehabilitation and training programs, understanding patients' limitations in visual function, and optimizing device designs and stimulation protocols. In the more mature field of cochlear implants, simulations have played a central role in optimizing electrode configurations and speech processing strategies (Hallum et al., 2007).

Early simulations relied on the view of phosphenes as the pixels of prosthetic vision. The perceived image is considered to be the grayscaled, downsampled version of the original. First attempts to mimic prosthetic vision used perforated masks of varying density and pixel count (Cha et al., 1992). Later, computer simulations and virtual reality environments (see for example Kasowski et al. (2021)) were used, allowing to include variable phosphenes size, brightness, and overlap (for a comprehensive review of phosphene-based models of prosthetic vision and their use in assessing functional capacity, see Chen et al. (2009a)). Chen et al. (2009b), for example, used such a simulation to design and assess rehabilitation programs for implant users. Functional consequences of image processing have been largely investigated using prosthetic vision simulators (see section 1.2.2). Denis et al. (2014) for example, evaluated an automatic text localization and specific rendering method to facilitate text detection. van Rheede et al. (2010) assessed the effects of full-field representation, zooming on regions of interest, and fisheye (high sampling frequency in the center, progressively lower resolution toward the edges) on visual performance. Wang et al. (2014) investigated the use of face detection-based image processing strategies to facilitate face recognition.

More recently, simulators were developed that were based on a more detailed biophysical description of the phenomenon (Golden et al., 2019; Beyeler et al., 2017a, 2019b). These advanced models (figure 1.13) highlighted the over-optimism of previous simulations. The *pulse2percept* model developed by Beyeler et al. (2017a) describes current diffusion in the retina, and the passing axons activation. The resulting tissue activation map is thresholded to predict RGCs activity. A series of temporal filters then describes the temporal evolution of the resulting percept. The model was fitted on patients' data (patients were asked to draw the percept evoked by electrodes activation), and its spatial predictions were validated. An interesting aspect of this modeling approach is that it combines knowledge of retina physiology and anatomy with results from psychophysics to predict the perceptual experience of patients. Importantly, this model allowed the authors to relate the variability between subjects to a few sets of parameters describing current diffusion and axon activation. The model of Beyeler et al. (2019b) was notably used, for example, to investigate the effects of perceptual learning (Wang et al., 2018) and image preprocessing (Han et al., 2021) on visual function.



Figure 1.13: Simulated prosthetic vision in Argus II users (adapted from Beyeler et al. (2017a, 2019b)). A. The patient is asked to draw the phosphene perceived upon single electrode activation on a touchscreen B. Example of drawings mapped on top of the corresponding electrode. Note the elongated shapes resulting from distal axons activation. C. Scoreboard model predictions for each electrode. The model assumes that each electrode elicits a circular phosphene; it is trained using all except the electrode of interest. D. Predictions from pulse2percept, which takes into account axons activation. The model is trained using all except the electrode of interest. E. Details of the pulse2percept model. The model takes into account current spread from the electrodes, as well as axon activation resulting in RGCs activation dependent on the fibers trajectories in the retina. The temporal sensitivity to current is modeled using an LNL filter. In particular, this filter takes into the effect of charge accumulation on neurons desensitization.

1.3 Methods for tuning visual implants

Much progress in artificial vision will come from improved stimulation protocols and image translation for the brain to make sense of the synthetic inputs it receives. However, precise stimulation of distinct RGC types is unachievable with current or in-development devices, making it impossible to replicate natural inputs and severely limiting current devices' performance. In the following section, I will describe the strategies that have been proposed to optimize how visual inputs are converted into stimulation patterns.

1.3.1 State of the art of visual implants tuning

Recordings-based strategies

Approaches to the optimization problem Most recent research on visual implant tuning follows biomimetic strategies, where the objective is to replicate the natural activity pattern of the visual system (Saal and Bensmaia, 2015; Choi et al., 2016; Valle et al., 2018). These methods require a good model of the retina neural code in order to define the target pattern of retina activity for any stimulus. Assuming this target pattern is defined, a first approach is to rationally design the stimulation protocol before interface use. However, this "open-loop" biomimetic strategy has strong limitations. Indeed, it requires building a model relating stimulation to retinal response, but because of the individualistic, dynamic and nonlinear nature of the inVisual stimulus

Target activation

Stimulation pattern



Recordings Retina

-

Tafazoli et al. (2020) used an algorithm analogous to simulated-annealing for adaptive closedloop stimulation of the visual cortex. After validation in a convolutional neural network (CNN) simulation, they applied their method in the visual cortex of mice. The mice underwent multisite electrical stimulation combined with multielectrode recording of neural activity. The authors demonstrated that their approach was successful in eliciting arbitrary-defined target responses in a population of neurons, despite noise and drift in the responses.

The main drawback of these approaches is that the whole optimization sequence would have to be repeated for any target optimization pattern. A solution is to learn a model of the retina response, then invert it for any new target pattern, without the need for new recordings.

• Learn then optimize

A promising approach to improve stimulation selectivity is to characterize the electrical receptive fields (eRF) of individual cells using reverse correlation techniques such as spike-triggered average (reviewed in Rathbun et al. (2018)), and use these eRFs to specifically target them. Briefly, the idea is to stimulate the retina with inputs sampled randomly from the space of electrical patterns (a requirement is that the distribution is spherically symmetric) and compute the average stimuli preceding a spike in the cell of interest over a predefined time window. The resulting stimuli average, called the eRF, is the linear filter of a linear-nonlinear model of the cell response behavior. An electrical stimulus that corresponds to this eRF maximally drives the cell.

Such LN models are good at describing RGCs behavior in response to direct multielectrode stimulation, but models involving more nonlinear interactions between electrodes are required to accurately describe network-mediated responses (Maturana et al., 2018).

Recently, Oesterle et al. (2020) proposed to infer the parameters of a very detailed biophysical model of the response of the retina to electrical simulation using a method for simulation-based Bayesian likelihood-free inference. In brief, the idea is to simulate the retina activity for candidate parameters, then compare the simulation results to actual data (two-photon measurements). Based on the discrepancy between simulation and data, new parameters are iteratively proposed until the model agrees with the data. The corresponding model was then used to find pulse shapes that would maximally differentiate OFF- vs ON-bipolar cells activation. In this paper, the authors focused on bipolar cells and used two-photon recordings, but this approach should generalize to RGCs with electrophysiological recordings.

• Hybrid approaches

Hybrid approaches use a combination of the two methods described above. Shah et al. (2019) proposed a closed-loop strategy in which an RGCs' responses dictionary is built by measuring the activity evoked by the 512 electrodes of their implant *in vitro*, with 40 different current amplitudes, 27 times each. Based on a linear model mapping ON and OFF parasol RGCs activity to visual percepts, the dictionary elements are combined in real-time using a fast and greedy optimization algorithm to minimize predicted percept distortion (measured using mean-squared error). This stimulation method is based on temporal multiplexing, the fact that sequential stimulation at a rate faster than the retina integration time leads to a unique percept. For a

1.3. Methods for tuning visual implants



Figure 1.15: Illustration of the principle of a retinal prosthetic encoder. A. We assume that to each visual stimulus corresponds an optimal electrical stimulation pattern. B. Most existing ACLS approaches infer, for a given visual stimulus, the corresponding best electrical pattern. However, this requires optimizing the electrical stimulation protocol for any new stimulus. C. Instead, we aim at learning the best mapping from visual stimuli to electrical patterns on a small set of stimuli, hoping that improvement will generalize to other stimuli. We term this mapping (in blue) the encoder. The problem is to find an encoder that matches the optimal encoder (in red). To do so, we parameterize the encoder and tune the parameters to maximize a measure of their quality (see chapter 2).

given visual stimulus, dictionary elements are sequentially combined based on the retina's response to the previous electrical stimuli in the sequence.

This strategy is hybrid in the sense that it learns a model used to define stimulation patterns, but then adapts the stimulation sequence online based on recorded activity. However, it is unclear whether this strategy could work with a more detailed nonlinear model of the mapping from evoked retinal activity to visual percepts. Indeed, the computation time for stimulus optimization needs to be shorter than the visual system integration time (tens of ms).

From pattern optimization to encoder optimization An issue with all the ACLS approaches mentioned above is that they consist in the direct optimization of the stimulation for a given activity pattern (blue arrow in 1.14). The optimization procedure would have to be repeated for any stimulus, which limits practical applicability and puts strong constraints on the computational efficiency of the algorithm used. We advocate for the use of a parametric transformation from images to stimulation patterns, that we call an encoder (Eckmiller et al., 1999). Instead of tuning the stimulation variables directly for each new stimulus, the parameters of the encoder would be tuned (red arrow in 1.14). This approach is depicted in figure 1.15). It essentially illustrates how the problem of finding appropriate stimulation protocols can be framed into a machine learning problem. Like any learning task, the encoder should lead to good performance on the training examples, but also have a performance that generalizes to other visual stimuli. To ensure generalization, the encoder should be optimized on a set of stimuli or tasks. Given the limited flexibility induced by the parameterization, the encoder cannot be optimal for every stimulus. As a consequence, the encoder optimization problem can be seen as a multiobjective optimization task. The encoder could be treated as a black-box or include detailed models of the visual system and electrode-tissue interface.

Defining the target pattern Even with a precisely controlled pattern of neuronal activation, the question remains of identifying appropriate target patterns. First, a model predicting the visual perception evoked by the retina activity is required. Even though such models exist for normal retina (see e.g. (Lindsey et al., 2019)), they may not be relevant for degenerated retina given the changes that occur during the diseases. Moreover, due to the imperfect selectivity of current devices, stimulation patterns that are suboptimal at replicating natural retinal activity may lead to better visual function. Consequently, the problem arises of defining a neural metric corresponding to the loss incurred by deviating from the target pattern. Tafazoli et al. (2020) for example, used the Euclidean distance between multi-unit activities, which seems arbitrary. Shah et al. (2017) addressed this problem in the case of RGCs stimulation by designing a method to learn a neural metric that gives a measure of the similarity of RGC population responses in a way that accurately reflects the similarity of visual inputs. This approach learns a metric by solving an optimization problem such that pairs of responses generated during different repeats of the same stimulus are closer than pairs of responses generated by different stimuli.

Another interesting approach to defining the target pattern was proposed in a recent paper by Romeni et al. (2021), who investigated a closed-loop optimization strategy *in silico* for an optic nerve prosthesis. For a given visual stimulus, either static or moving, a visual cortex target activation pattern is computed based on a CNN model of the visual pathway. This CNN model was previously trained to identify hand-written digits and objects on standard image datasets (MNIST and FMNIST). The assumption is that by evoking a cortical activity closed to the one occurring in natural vision, the perceptual experience will be more natural and qualitative. This idea is supported by recent developments demonstrating that CNNs could be used as accurate models of the visual system (Kriegeskorte, 2015; Yamins and DiCarlo, 2016; Lindsey et al., 2019). The activation of the visual system upon optic nerve stimulation is modeled using a simple model of the nerve stimulation and the CNN model of the visual system, and an optimization procedure is then run to find the stimulation pattern that evokes a cortical activation as close as possible to the target pattern.

The authors then assessed the accuracy of the CNN when classifying images based on artificial inputs. The evolutionary heuristic led to the dramatic improvement in the ability of the model to identify stimuli. The authors could also investigate the optimization performance depending on the target visual pathway stage (the one that is recorded) and the simulation accuracy. However, the authors assumed perfect recordings of the target cortical layer activity and did not investigate the consequences of imperfect measurements. Moreover, they assumed perfect matching between the model used to determine the target activation pattern and the visual system. Finally, their optimization strategy was designed to optimize the stimulation on a very restricted set of stimuli, and transfer of improvement to other stimuli was not assessed. In general, an issue with the use of such detailed surrogate models of the visual system is the fact that they are not patient-specific. Given the significant changes that occur in the retina and the cortex following degeneration (see 1.1.2 and 1.1.2), it is not clear whether existing models will be adequate.

In general, a strong limitation of the recordings-based strategies is that they define the target pattern of activity without considering changes that occur in downstream visual areas as the disease progresses and the visual system adapts to prosthetic stimuli. They all assume a static ideal decoder able to optimally interpret the outputs of the retina. As a consequence, even if we had a perfect model of natural visual processing, it may lead to selecting target activity patterns that are suboptimal. This suggests that the target activity pattern should itself be adapted in a closed-loop.

In summary, despite the attractive possibility to mimic natural RGCs activity, recordingsbased strategies still face numerous limitations that make them out of reach with current technologies: the requirements for good models of visual processing, the ability to define optimal target activity patterns, and the ability to stimulate and record at the same time.

Perception-based strategies

Given the numerous challenges that recording-based strategies entail, another avenue to ACLS is to directly measure the perception evoked by electrical stimulation. Contrary to recording-based strategies, this approach does not require defining the target retina activity pattern or modeling the coupling between implant and retina activity. Moreover, it can directly be applied to existing devices, given that it does not require modifying the hardware.

In general, a perception-based optimization strategy implies several challenges. First, it requires a relevant metric to quantify visual performance or the quality of patients' visual experience. This is difficult, as a patient may perform well in a very specific task (such as face recognition) with a given encoder, without being able to perform other, as important tasks (such as spatial navigation).

Moreover, the perceptual experience with a given encoder may lead to good visual performance, while being exhausting, making the encoder impractical. Second, subjective reports may be unreliable and are typically slow to gather. As a consequence, the data needed to tune the encoder may require days of measurements. The type of feedback queried from patients is likely to be paramount to efficiently optimize the system. For example, it has been proposed that the measurement of the oculomotor behavior, and in particular the microsaccade pattern, could be used as an objective measure of vision and could help fine-tune the stimulation parameters (Stingl et al., 2013). This is based on the idea that higher quality vision would induce more natural oculomotor behavior. However, this method would only apply to devices that preserve free eye movement.

Eckmiller et al. (1999, 2005, 2007) introduced such a perception-based closed-loop optimization strategy and assessed it on healthy subjects using simulated prosthetic vision. Subjects had to compare the perception elicited by a given stimulation to an expected percept, such as a white circle on a dark background, and new parameters were iteratively proposed based on reinforcement learning and genetic algorithms. After 60 to 120 minutes, the algorithm converged towards the encoder for which the percept looked most similar to the shape. This method, however, was not demonstrated on patients, nor with a validated prosthetic vision model. Moreover, it is based on the perceptual comparison between an input pattern and a perceived pattern, but patients do not have access to the input pattern. Lastly, Eckmiller (1997); Eckmiller et al. (1999, 2005) did not show that the optimization procedure led to improvement in sight restoration, and the authors do not mention whether the stimulation parameters optimized for one input pattern were also optimal for other patterns.

1.3.2 State of the art in human-in-the-loop optimization

A dialogue between human and machine

In the previous section, I described the solutions that were proposed to the problem of retina implant tuning. The approach we will focus on in the remaining of this thesis is to use a parametric encoder converting images to electrical stimulation patterns. The problem is to select the parameters that maximize some measure of visual function. In fact, retinal prostheses are one of many occurrences of a human-machine interface that requires a precise although extremely challenging tuning through human interaction. This is denoted as "human-in-the-loop optimization" (HILO). Typical examples that attracted much research are A/B testing for user interfaces or recommender systems (Lü et al., 2012). In all these problems, a human user interacts with a system and provides a measure of performance (such as sales, click-through rates, or energy expenditure) according to which the system should be optimized. We can draw from successes in this broader field for our specific application.

These problems all share the property that the objective function is a black-box, i.e., its analytic description is unknown, and we do not have access to derivative information, making it difficult to optimize. Moreover, in a subset of these problems, function evaluation is expensive in terms of money, time, or any other resource. When humans enter the loop, additional issues arise, such as exertion, fatigue, or even learning. This typically limits the amount of data we can gather from humans. The limitation on the amount of data we can collect makes the optimization more sensitive to noise that arises, for example, because of the unreliability of subjects' responses.

Basic methods such as random search, grid search, or discrete steps find optimal configuration at a high cost and are limited to low-dimensional problems (Bergstra and Bengio, 2012). As a consequence, a central issue is to select new candidate parameters that are informative about the best configuration. However, compared to other expensive optimization problems where the computation time to select new samples is usually not an issue, it can become a critical aspect of HILO.

HILO methods can be subdivided depending on the type of feedback. It can either be based on a direct measure of the performance of the system, typically when a performance metric is accessible or when the system either passes or fails a test; or be based on comparisons between states. I term these two distinct scenarios performance-based and preference-based optimization. In the case of visual prosthetics, a performance measure would correspond for example to visual acuity, whereas a preference measure would be the relative quality assessment of two percepts evoked by two distinct encoders for the same visual stimuli.

Performance-based optimization

In most applications of HILO, an objective measure of the system performance when interacting with humans is defined and optimized. For example, Zhang et al. (2017) applied Covariance Matrix Adaptation Evolution Strategy (CMA-ES) to the optimization of four control parameters of an ankle exoskeleton. In that case, the objective was energy expenditure when walking at constant speed on a treadmill, measured through dioxygen consumption and carbon dioxide production. The optimization process resulted in a significant reduction in energy 1.3. Methods for tuning visual implants



Figure 1.16: Principle of Human-in-the-loop optimization. A measure of the quality of the interaction between the machine and the user is recorded. Accumulated data are used in conjunction with prior information to propose new control parameters for the control architecture of the machine.

consumption, 5.8 $\% \pm 6.2\%$ lower than hand-tuned parameters. In addition, the authors noted that in pilot tests, open-loop model-based optimization techniques were ineffective because of sensitivity to noise and adaptation dynamics. These results suggest that human-in-the-loop methods are more robust than model-based methods.

Recently, Welker et al. (2021) applied CMA-ES to tune the control parameters of an anklefoot prosthesis used by patients with unilateral transtibial amputation. Four different controller architectures were tested, with 2 to 5 parameters. The authors failed to observe a significant reduction in energy expenditure compared to generic parameters, highlighting the difficulty of designing an effective HILO procedure. Possible reasons could be the choice of the optimization algorithm, objective function, motor adaptation, and learning.

Another related example is the application of Bayesian optimization (described in section 3) to the problem of tuning two control parameters of a soft exosuit for hip assistance during walking (Ding et al., 2018). Although the optimization led to lowered metabolic expenditure, the authors did not compare the performance of this algorithm to other methods.

Bayesian optimization has also been applied to Deep Brain Stimulation (DBS). In DBS, settings of the stimulation protocol are considered the most critical factors for functional outcomes, along with accurate electrode placement (Hell et al., 2019). However, tuning is done manually by a clinician in a time-consuming trial and error procedure. This method precludes from using potentially more efficient advanced combinations of parameters, and, as a consequence, ACLS methods are considered very promising (reviewed in Parastarfeizabadi and Kouzani (2017)). Grado et al. (2018) showed in a computational model of Parkinson's disease that Bayesian optimization could be used to efficiently reduce the power of beta oscillations in the basal ganglia, which is associated with movement disorders. Bayesian optimization was more efficient and more robust to noise than the Nelder-Mead and DIRect algorithms, two popular methods for non-convex global optimization.

Very recently, Bayesian optimization was applied to the problem of tuning the frequency of

deep brain stimulators in two patients with Parkinson's disease (Louie et al., 2021). The authors optimized a measure of forearm rigidity. However, in this application, the advantage of using Bayesian optimization compared to other methods such as random sampling is not clear. This may be since the problem is one-dimensional and so too simple to see a noticeable difference.

Reinforcement learning was also applied in HILO (Sutton and Barto, 2018; DiGiovanna et al., 2009). Thomas et al. (2009) applied Actor-Critic Reinforcement Learning (ACLR) to control functional electrical stimulation of a human arm in simulations. However, ACRL approaches usually rely on predefined, a priori reward signals based on expert knowledge. To address this, Pilarski et al. (2011) used ACLR to allow subjects to optimize the policy of an electromyography-controlled prosthesis arm via their positive and negative feedback signals. The system was able to learn a two-joint velocity control task under continued user input without the need to integrate detailed tasks and domain knowledge into the learning framework.

Preference-based optimization

In many situations, it is challenging to find a metric to quantify the performance of a system, and humans are better at evaluating differences rather than absolute magnitude (Kahneman and Tversky, 1979). For example, intuitively, it is much easier to compare pairs of movies than to rate them on a scale from 0 to 100. Such absolute rating may be noisier and lead to inconsistent results due, in particular, to the order of presentation (Yannakakis and Hallam, 2011).

As a consequence, preferences tend to be more reliable than numerical scores, as was shown in domains such as information retrieval (Chapelle et al., 2012) and autonomous driving (Basu and Ghosh, 2017). Moreover, in many situations, several objectives have to be jointly taken into account, and a choice has to be made to select the point of the Pareto front (the set of points for which no objective can be improved without making another worse) to report as a solution. In the field of engineering, preference elicitation has been used to solve the problem of balancing the different objectives of multi-objective optimization problems (Lepird et al., 2015).

For example, when tuning the frequency of deep brain stimulators, side effects such as bradykinesia and tremor can appear due to current spread that are difficult to quantify. For this reason, Louie et al. (2021) measured the preference of patients for pairs of stimulation frequencies and showed that the preferred parameters were not those minimizing forearm rigidity. They suggest that preference data could be used as a way to optimize deep brain stimulation protocols in the future.

The idea that preference measures may be more relevant than objective performance is further supported by the work of Tucker et al. (2019), who applied Preference-based Bayesian optimization (PBO) to the problem of optimizing two gait parameters of the Atalante lower-body exoskeleton. The goal was to maximize user comfort when walking inside the exoskeleton. With each new set of parameters, the subject had to compare the current set of parameters with the previous one for 20 iterations. At each iteration, new parameters were selected adaptively so as to efficiently find the preferred parameters. Interestingly, no correlation was found between metabolic expenditure and user preferences. In this study, the authors included coactive feedback, where the subject could suggest modifications to the parameters (such as 'slightly shorter steps'), which improved overall performance.

1.4. Thesis outline

PBO was first applied to procedural animation design (Brochu et al., 2010a). Later, it was used to tune hearing aids (Nielsen, 2015). In the application to hearing aid, the authors noticed that listener fatigue and training effects appeared to noticeably influence the consistency of subjects. However, the algorithm proved to be robust to this noise. Moreover, the stimulus (music) was kept constant during the optimization procedure, and the authors did not investigate if the preferred settings could generalize to other stimuli. This highlights a general problem of HILO methods: since the optimization procedure is run in a limited set of conditions, the observed improvements may not generalize. The choice of the objective is likely to be paramount for generalization. To solve this problem of generalizing improvement across many conditions, Kupcsik et al. (2018) applied PBO to optimize a robot-to-human object handover across a variety of randomly chosen contexts (objects size and type) based on human feedback.

To our knowledge, the first application of PBO to a brain-computer interface was reported by Lorenz et al. (2019), who used it to maximize the phosphenes brightness elicited in transcranial alternating current stimulation when tuning current frequency and phase.

However, in none of the aforementioned studies did the authors compare their acquisition rule (the criterion used to select new samples) to, for example, random sampling. It is thus impossible to conclude whether the use of such a complex algorithm was needed compared to mere preference learning.

From these examples, we can draw a few conclusions. First, to compare results from different HILO studies and to validate the use of advanced algorithms, control experiments in which random acquisition is performed are required. If possible, a comparison with generic or handtuned parameters should also be performed. Moreover, when the system can work in various conditions, generalization of improvements to other contexts than those of the optimization runs should be performed.

An illustration of the application of HILO to the problem of tuning a retina implant encoder is presented in figure 1.17.

1.4 Thesis outline

In this thesis, I will introduce a closed-loop approach in which subjects' responses in a visual task would be used to continuously optimize the way the implant converts visual inputs into stimulation patterns. The problem of optimizing retinal implant encoders can be decomposed into three parts :

- First, we need to find a good encoder parameterization, that is complex enough to capture the complexity of the implant-retina interface and simple enough to be optimized.
- Second, we need to define a measure of the encoder quality, that can be used as an optimization metric.
- Third, we need to find an optimization algorithm that is efficient in solving our problem given the type of data we can expect from patients.

To determine whether our strategy is successful, I set the following criteria :



Figure 1.17: Retinal prosthetic encoder optimization. A. The goal is to find a retina implant encoder (in blue) that leads to a perceptual quality closed to the one obtained with an unknown optimal encoder (in red). B. Different encoder parameters lead to visual experiences of various qualities. By definition, the optimal encoder (red) corresponds to the highest perceptual quality for a given stimulus. The goal is to adjust the parameters of our candidate encoder to maximize the quality of the visual experience. C. In performance-based optimization, we assume that an encoder's quality is related to the performance in a visual task. The best parameters are found by maximizing the performance. D. In preference-based optimization, we assume that an encoder's quality is related to a latent value function that we can only evaluate through qualitative comparisons. The best parameters are found by maximizing the latent value.

- Given a realistic data budget, the optimization strategy should find an encoder that consistently leads to a higher optimization metric value compared to controls to be defined.
- The performance improvement should generalize to visual stimuli and tasks that are distinct from the ones used during the optimization process.

In chapter 2, I develop solutions to these different problems and assess the criteria above. In particular, I investigate the potential of preference comparisons for different encoders, as well as a performance measure in visual acuity-related tasks. The needs of this application led me to highlight the limitations of optimization algorithms in use nowadays. In chapter 2, I present an in-depth review of these optimization algorithms, which is currently missing, and propose modifications that lead to significant performance improvement. These developments, from theoretical aspects to practical implementation considerations, were critical in the success of the work described in chapter 1.

The generalization of improvement from one visual task to another is a tremendous challenge. A solution is to optimize the encoder across various tasks. However, all may not be equally informative about the best configuration. For example, in performance-based opti-

1.4. Thesis outline

mization, if the task is too easy or too difficult, responses will all be equivalent, whatever the encoder. The problem of optimizing a system working in various contexts is not new, but to our knowledge, the idea of joint optimization and adaptive context selection to ensure informative outputs has never been considered. In chapter 3, I will analyze this general problem and propose efficient algorithms to tackle it. In particular, I apply this new framework in the context of visual psychophysics.

2

Human-in-the-loop optimization of visual prosthetic stimulation

2.1 Introduction

Retinal prosthetics aim at restoring vision by artificially stimulating neurons in the retina (Lorach et al., 2013). However, currently, the visual resolution that can be achieved with such implants is very low (Stingl et al., 2017; Palanker et al., 2020). This is due in particular to the distortion incurred by the small number of electrodes, the activation of passing axons, and the unspecific stimulation of cells of different types (Golden et al., 2019; Beyeler, 2018; Beyeler et al., 2019b). Appropriate encoding of visual inputs into electrical stimulation patterns has the potential to greatly reduce these factors, and improve visual function (Spencer et al., 2018; Shah et al., 2019). Here, we asked whether it is possible to improve the performance of these devices by using patients' feedback in visual tasks to optimize the stimulation protocol.

2.1.1 Optimization of retinal prosthetic encoders

In existing retinal prosthetic devices, the patient wears a glasses-mounted camera that captures the visual scene. This visual input is processed by a small pocket computer (the visual processing unit) and converted into electrical pulses through a mathematical function that we call an encoder. The visual implant's processing unit gives the opportunity to use image processing algorithms to adapt the electrode array activity (or the light pattern in the case of microphotodiodes arrays).

Much research effort has been devoted to optimizing the mapping from visual stimuli to electrical stimulation patterns. On the computational side, two main strategies have been pursued: the first is to try to replicate the inputs that RGCs receive in normal conditions (Nirenberg and Pandarinath, 2012; Shah et al., 2019). At a cellular level, for example, much work has been done to perform specific stimulation of ON and OFF RGCs (Twyford et al., 2014; Guo et al., 2018; Lee and Im, 2019). However, applying these methods in real-life patients requires the ability to record the elicited neuronal activity and identify RGC types in patients' retinas, and thus significant technological developments. Moreover, we currently do not have a good model to predict how visual inputs are transformed into electrical signals in the retina, and even if such a model were available, the limited resolution and signal-to-noise ratio that can be achieved would make it difficult to match the stimulation that RGCs receive under natural conditions.

The second strategy is to adapt the transformation from visual inputs to electrical stimulation (the encoder), to directly improve the patient's visual performance, for example, by highlighting specific features of the visual stimuli such as edges or motion (Feng and McCarthy, 2013; Luo and da Cruz, 2016; Barnes et al., 2016). However, previous research has shown that the perceptual experience of patients equipped with retinal prosthetic devices varies significantly from one patient to another (Beyeler et al., 2019b; Erickson-Davis and Korzybska, 2020), suggesting that stimulation protocols should be carefully adapted in a personalized way, for example by using feedback from patients. Moreover, because of the limited bandwidth of these devices, it is not clear which features of the visual scene should be emphasized to improve visual function.

2.1.2 Perception-based optimization of retina implant encoder

Here, we follow the second approach and propose a new, readily applicable strategy based on machine learning to optimize how visual signals are encoded into stimulation patterns. Our approach is based on a model of prosthetic vision predicting visual percepts from electrodes activation (Beyeler et al., 2019b). We show that, if we know this model, it is possible to compute the 'optimal' encoder, a mapping from visual inputs to stimulation patterns, that minimizes percepts distortion. This encoder is determined by parameters that can be tuned in a subject-specific manner. However, the model may not be a perfect description of real-life patients' perceptual experience, and the perceptual model parameters can only be inferred with limited precision. Moreover, implant users emphasize the fact that prosthetic vision is fundamentally different from natural vision, and the experience of prosthetic vision is highly subjective and variable (Erickson-Davis and Korzybska, 2020). As a consequence, and given the limited bandwidth of these prosthetic devices, minimizing percept distortion may not be the best target to recover useful vision.

To address these issues, we introduce a closed-loop optimization in which subjects' responses in visual tasks are used to continuously improve the way the implant converts visual inputs into stimulation patterns. To take the subjective aspects of prosthetic vision into account, the task we use to optimize the encoder consists of pairwise comparisons between candidate encoders. As this approach requires no new hardware, it should be relatively easy to implement with existing technologies.

Given the limited amount of data we can gather from individual subjects, sample-efficiency is paramount. Moreover, these data can be relatively unreliable. Our method is thus based on a state-of-the-art preferential Bayesian optimization algorithm. To overcome some limitations of existing algorithms, we proposed new implementation methods and a new sample acquisition rule that outperforms existing methods (we will consider these aspects in more detail in chapter 3). We demonstrate on a validated prosthetic vision simulation that our method leads to significant and robust improvements in perceived images quality and visual acuity.

2.2 Framework

For simplicity, let us consider a static image s (vector of size $n_s \times 1$, where n_s is the number of recorded image pixels) which results in a pattern of electrodes amplitudes a (a_i is the current amplitude at electrode i, a is a vector of size $n_e \times 1$, where n_e is the number of electrodes). We define a retinal prosthetic encoder as the mapping from input stimulus s to amplitudes a. For

a linear encoder *W* (matrix of size $n_e \times n_s$), we have a = Ws. Our goal is to find an encoder *W* that maximizes the perceptual quality experienced by the patient for any input *s*.

We consider that the only available data are subjects' perceptual reports, and we aim at learning the best encoder by measuring the subject's responses in visual tasks with different encoders. The main challenges we faced are the potentially high dimension of the encoder, the unreliability of subjects' responses, and the very limited amount of data we can collect. For simplicity, we do not take into account the temporal dimension: we consider a static model, where the electrical pulse shapes and frequency are kept fixed, and only the pulse amplitudes vary, but our approach could generalize to include these variables.

2.2.1 Strategy assessment on normally sighted subjects

In order to develop our approach, we experimented on healthy subjects using the prosthetic vision simulator *pulse2percept* (Beyeler et al., 2017a) (figure 2.1.A). This model simulates the visual percepts experienced by patients fitted with a prosthetic device. In particular, it describes the effects of current diffusion and passing axons activation on the perceived phosphenes. It was fitted and validated using data from users of the epiretinal implants Argus I and Argus II (Second 79 Sight Medical Products, Sylmar, CA) (Horsager et al., 2009, 2010, 2011; Nanduri et al., 2012; Beyeler et al., 2019b). Here, we focused on Argus II, a 60-electrode epiretinal device that has been implanted in approximately 300 individuals since its commercial approval in the EU in 2011, the US in 2013, and Canada in 2015 (Ahuja et al., 2011; Humayun et al., 2012; Rizzo, 2011; da Cruz et al., 2016).

The *pulse2percept* model assumes that electrodes currents combine linearly to produce a spatial intensity profile I(x, y), where (x, y) are spatial coordinates on the retina (Beyeler et al., 2019b). This profile is then transformed to take into account passing axons activation to give an activation map $I_{axon}(x, y)$ which, again, combines electrodes' contributions linearly. This profile is then passed through a nonlinearity so that the static version of *pulse2percept* is essentially a linear-non-linear model.

Let ϕ denotes the parameters describing the *pulse2percept* model for a specific subject. The prediction of the model for any pattern of electrodes amplitudes a is: $p_{\phi}(a) = \min(|M_{\phi}a|, m)$, where p is the subject's percept (vector of size $n_p \times 1$, where n_p is the number of pixels in the percept), m the maximal percept luminance (m = 1) and M_{ϕ} the matrix of size $n_p \times n_e$ whose rows correspond to electrodes projective fields. Since this formulation of the *pulse2percept* model is not apparent in its original publication, we compared the percepts predicted by the original software to those predicted with our simpler linear-nonlinear formulation. We generated 10^4 random amplitude patterns and computed the difference in the predictions of the two versions of the model, keeping all parameters of the *pulse2percept* model at their default values. The root-mean-squared error was 2.8×10^{-5} , meaning that the two versions indeed correspond to the same model.

2.2.2 Encoder parameterization

The encoder matrix W has size $n_e \times n_s$. Learning these $n_e \times n_s$ elements based on subjects' responses seems infeasible. Therefore, to reduce the dimensionality of the optimization prob-



Figure 2.1: A. Encoder parameterization. Eight parameters ϕ specify the perceptual model described by \mathbf{M}_{ϕ} : the current diffusion coefficient, axon sensitivity to currents, two axon shape descriptors, implant position and orientation, and retina sensitivity to current. Columns of \mathbf{M}_{ϕ} correspond to each electrode projective field, i.e. the linear contribution of a given electrode to each percept's pixel brightness. Axon activation induces elongated phosphenes along the fibers. The optimal linear code $\mathbf{W}_{\phi} = \mathbf{M}_{\phi}^{\dagger}$ corresponding to the linear approximation of the perceptual model is the pseudo-inverse of \mathbf{M}_{ϕ} . On the contrary, a naive linear code simply downsamples the input image to the number of electrodes. The encoders are depicted by plotting the receptive field of an electrode of the implant. Note the center-surround structure of the optimal receptive field. To each encoder corresponds a different pattern of current amplitudes. The pulse2percept model predicts percepts that are much less distorted when using matrix $\mathbf{M}_{\phi}^{\dagger}$ to encode the input signal. **B**. Visual preference task. A visual stimulus **s** is converted with two different encoders ϕ and ϕ' into stimulation patterns $a(\phi, \mathbf{s})$ and $a(\phi', \mathbf{s})$. The corresponding percepts are simulated using the pulse2percept model (determined by a subject's parameter ϕ^{\star} , and displayed on a monitor. The subject is asked which percepts she finds less distorted, and response z is recorded. New duels are proposed and after 60 iterations, the best encoder is inferred based on the subject's preference.

lem, we parameterized the encoder by making use of the *pulse2percept* model (Beyeler et al., 2017a). This can be seen as a way to integrate prior knowledge from biophysics into the optimization problem. Eight patient-specific parameters ϕ were used to parameterize the encoder: the current diffusion coefficient, the axon sensitivity to current, the implant location and orientation on the retina, and two parameters that describe the fibers bundles trajectory in the retina (Jansonius et al., 2009). The last parameter was used to scale the current intensity at each electrode. Based on these parameters, *pulse2percept* predicts the projective field of each electrode in M_{ϕ} .

Given M_{ϕ} , we compute the linear encoder W_{ϕ} that would minimize the signal distortion defined as the euclidean distance between the input image *s* and the predicted percept *p*, assuming that projective fields combine linearly (that is, removing the modulus):

$$W_{\phi} = \underset{W}{\arg\min} \left\langle (s - M_{\phi} W s)^2 \right\rangle$$
(2.1)

The optimal linear encoder W_{ϕ} is simply the pseudo-inverse of M_{ϕ} : $W_{\phi} = M_{\phi}^{\dagger}$ (figure 2.1.A). That way, the encoder W_{ϕ} is entirely determined by the eight parameters ϕ . The rows of W_{ϕ} can be interpreted as each electrode receptive field, i.e., how each pixel in the input image participates in the activity of the electrode.

Ideally, we would instead solve the exact minimization problem:

$$\underset{W}{\arg\min} \langle (s - \min(|M_{\phi}Ws|, m))^2 \rangle, \qquad (2.2)$$

but this would require performing a slower optimization (for example using quadratic optimization methods). Removing the modulus prevents from finding solutions where electrodes currents could be combined to generate a negative current inducing retinal activation, but simplifies the computation of the encoder. In practice, numerical simulations showed that there was virtually no difference between percepts obtained using the encoder $W_{\phi} = M_{\phi}^{\dagger}$, and those obtained with an encoder that was the numerical solution of 2.2 (see supplementary figure 2.10 for an example).

This parameterization of the encoder can be seen as a global shaping of the percepts. Traditional approaches combine elementary phosphenes to create a meaningful image, and electrodes crosstalk and current spread are considered detrimental. With the method we propose, electrodes interactions are leveraged to shape the percept globally.

2.2.3 Motivation of the closed-loop optimization approach

For a patient whose mapping from electrical pulses to percepts can be described with parameters ϕ^* , assuming the *pulse2percept* model is a good description of her prosthetic vision, we thus have a way to compute the encoder minimizing perceptual distortion on average. Figure 2.1.A illustrates the dramatic improvement that such an "optimal" encoder brings to perceptual quality. Here we see that a 'naive encoder' where each electrode responds to visual pixels in a small region of space results in a highly blurred visual percept. In contrast, optimizing the projective field of each electrode can act to compensate for current spread and axon activation. On figure 2.1.A, the optimal linear receptive field for a single electrode is plotted. This receptive field displays a center-surround structure where the electrode is driven by bright pixels in the center and inhibited by bright pixels in the surround. This center-surround structure is responsible for confining the region of activation, hence the resulting phosphenes.

However, in practice, the *pulse2percept* model may not be a perfect description of subjects' experience, and the parameters ϕ^* are known with limited precision. Moreover, they vary widely from patient to patient (Beyeler et al., 2019b; Jansonius et al., 2009). Further, it is not clear which metric could be used as a good proxy for prosthetic visual function. Whereas metrics exist to quantify image quality as experienced by users with normal vision (Simoncelli et al., 2004), the fact that patients report the experience of prosthetic vision as fundamentally different from natural vision suggests that such metrics may be irrelevant (Erickson-Davis and Korzybska, 2020). As a consequence, the open-loop approach which consists in inverting the perceptual model p_{ϕ^*} may likely lead to suboptimal encoders.

For this reason, we propose to adjust the parameters ϕ by using a closed-loop optimization approach based on subjects' responses in visual tasks. Our optimization problem is to find the parameters ϕ of the encoder W_{ϕ} that maximizes some measure of visual function $g(\phi)$:

$$\underset{\phi}{\arg\max} g(\phi) \tag{2.3}$$

Of course, *g* is not known to us a priori and must be inferred using subjects' feedback. In the following section, we describe a way to do this, by asking patients to perform pairwise comparisons between pairs of encoders (ϕ_1 , ϕ_2).

2.2.4 Preferential Bayesian Optimization

One challenge we faced was to design a visual task based on which we could optimize the encoder. We sought a task that would provide a broad measure of how 'good' subjects' visual percepts are for a given ϕ , and avoid focusing on a specific visual feature that was useful for performing the task at hand. To do this, we focused on a type of task called 'preference learning', where one iteratively ask the subject, for a set of stimuli, which of two encoders leads to the highest perceptual quality, i.e., which of the two leads to a percept that resembles the described stimulus the most, and use preference-based optimization to infer the preferred encoder (figure 2.1).

This human-in-the-loop optimization strategy comes with several challenges: the limited amount of data we can collect from patients and the fact that these data are unreliable. Indeed, the response timing in visual tasks for prosthetic device users can be very slow. Barnes et al. (2016), for example, measured average response timings in a light-localization and a grating acuity task of the order of 50s for suprachoroidal implant users. This is most likely because of the very restricted field of view which requires prolonged head-scanning. Further, responses can be inconsistent due, e.g., to fatigue or attentional lapses, but also because of the trial-to-trial variability of percepts evoked by a given electrical stimulation pattern (Beyeler et al., 2019b).

To address these issues, we use Preferential Bayesian Optimization (Brochu et al., 2010b; Dewancker et al., 2017; Gonzalez et al., 2017). Bayesian optimization is a powerful method for derivative-free optimization of black-box functions, that is, functions whose analytical form is unknown, and for which no derivative information is available (Frazier, 2018; Brochu et al., 2010b; Shahriari et al., 2016). It is especially relevant when function evaluation is costly, such as in the case we are interested in, as human feedback takes valuable time. In specific situations, direct evaluation of the objective function is impossible: we can only evaluate order relationships between pairs of outputs in so-called "duels". This arises, for example, when trying to optimize a system according to the preferences of people, and is thus referred to as Preferential Bayesian Optimization (PBO) (Gonzalez et al., 2017) or Preference-based Bayesian Optimization (Dewancker et al., 2017). Preference-based optimization is especially useful when the optimization process is based on human feedback. Indeed, for humans, comparison is much easier than rating (Kahneman and Tversky, 1979; Shah et al., 2014), and in many situations, the value of the function to optimize is not directly measurable (Brochu et al., 2010b,a). PBO has been successfully applied to the problem of tuning hearing aids (Nielsen et al., 2014), where participants gave their preference between combinations of two parameters when listening to music, and to transcranial alternating current stimulation (Lorenz et al., 2019), where the objective was to maximize the perceived phosphene brightness by tuning two stimulation parameters.

Classical Bayesian optimization (BO) proceeds in two steps: first, a surrogate Bayesian model of the objective function is learned using previous function evaluations, typically using Gaussian process regression (Rasmussen and Williams, 2006). Then, a new sample is adaptively chosen according to a predefined heuristic implemented in an acquisition rule. This heuristics is used to tradeoff exploration of the parameters space (selecting samples in regions with high uncertainty about the model) and exploitation (sampling in regions where the model predicts with high confidence that the encoder is good). The main difference with PBO is that, in the latter, the acquisition function returns a new duel instead of a single input.

2.3 Methods

We presented simulated percepts generated by the *pulse2percept* model on a monitor at eyes height, approximately 50 cm away from the participants. 12 out of 24 subjects participated in the experiment remotely thanks to the remote access software TeamViewer (TeamViewer GmbH). We did not observe differences in the results between remote participants and participants who performed the tasks in the lab (supplementary figure 2.11). The instructions were displayed on the monitor, and the participant answered using the computer keyboard. All participants had normal or corrected to normal vision. One of them had prior experience in psychophysics experiments. Tasks were assigned in random order to avoid the effects of fatigue or attention lapses. During preference comparisons, the input stimuli were letters that covered 19° of the visual field.

2.3.1 Principle of Preferential Bayesian optimization

In PBO, we assume that each input x from the search space \mathcal{X} has an associated value (in our application, this value represents the implicit value of a given encoder), determined by a real valued function g, and the outcome of any duel (x, x') is determined by the values difference f(x, x') = g(x) - g(x'). We are therefore interested in solving the global optimization problem:

$$\boldsymbol{x}^{\star} = \underset{\boldsymbol{x} \in \mathcal{X}}{\arg\max} g(\boldsymbol{x}) \tag{2.4}$$

Here, we consider that we measure binary outcomes $x \succ x'$, or $x \prec x'$. We have access to previous duel results \mathscr{D} and we assume that we will perform *T* duel evaluations before having to report a solution to 3.1.

We model the preference, noted *y*, as a stochastic process, such that:

$$\mathbb{P}(y=1|x,x',f) = \mathbb{P}(x \succ x'|f) = \Phi(f(x,x')),$$
(2.5)

where Φ is the Gaussian cumulative density function.

Houlsby et al. (2012), based on previous work on preference learning with Gaussian processes (Chu and Ghahramani, 2005), proposed to model preference learning as a special case of Gaussian process binary classification (Rasmussen and Williams, 2006). The latent function g is modeled using a Gaussian process (GP) with kernel k and mean m. This induces a GP prior $f \sim \mathcal{GP}(m_{\text{pref}}(\cdot), k_{\text{pref}}(\cdot, \cdot))$ on the preference function f, with prior mean $m_{\text{pref}} : \mathcal{X}^2 \to \mathbb{R}$ and a preference kernel $k_{\text{pref}} : \mathcal{X}^2 \times \mathcal{X}^2 \to \mathbb{R}$, defined as:

$$k_{\text{pref}}((x_i, x_j), (x_k, x_l)) = k(x_i, x_k) + k(x_j, x_l) - k(x_i, x_l) - k(x_j, x_k)$$
(2.6)

The posterior distribution is approximated using Expectation Propagation (Minka, 2001; Seeger, 2002).

After spending the sampling budget, the reported solution to problem 2.4 is the maximum x^* of the posterior mean of the Gaussian process $p(g|\mathscr{D})$. This corresponds to a risk-neutrality assumption (Berger, 1985), i.e. we value x^* according to its expected value.

2.3.2 Maximally uncertain challenge acquisition

Inspired by the challenge acquisition rules used in dueling bandits (Bengs et al., 2021), we propose a new acquisition function for preference-based optimization (I describe in detail the motivations for this acquisition rule in chapter 3). The first duel member x is chosen to be the maximum of the Gaussian process posterior mean. The second duel member x' is the challenger for which the epistemic uncertainty of the outcome is the greatest, i.e. the maximum of $\mathbb{V}[\Phi(f)]$ (see 3.3.2). This variance can be computed analytically (see appendix D) to give:

$$\mathbb{V}\left[\Phi\left(f(\boldsymbol{x})\right)\right] = \mathbb{E}\left[\Phi\left(f(\boldsymbol{x})\right)\right] - 2T\left(\frac{\mu_f(\boldsymbol{x})}{\sqrt{1 + \sigma_f^2(\boldsymbol{x})}}, \frac{1}{\sqrt{1 + 2\sigma_f^2(\boldsymbol{x})}}\right) - \mathbb{E}\left[\Phi\left(f(\boldsymbol{x})\right)\right]^2, \quad (2.7)$$

with the Owen's T function defined as:

$$T(h,a) = \frac{1}{2\pi} \int_0^a \frac{e^{-\frac{1}{2}h^2(1+x^2)}}{1+x^2} dx$$
(2.8)

The Owen's T function can be cheaply evaluated numerically Young and Minder (1974); Thomas (1986). The computation of the gradient of $\mathbb{V}[\Phi(f(x))]$ with respect to x is thus straightforward (see appendix D).

We evaluated this acquisition function on a set of 34 synthetic problems and showed that it outperforms state-of-the-art acquisition functions (see 3.4.2).

2.3.3 Leveraging knowledge from previous patients.

Constraints on the search space

The choice of the search space for the encoder parameters is paramount. Indeed, the wider the search space, the more difficult the optimization, but on the other hand, it should be large enough to include the global optimum. To determine the range of values for the current diffusion and the axon sensitivity to current, we made use of the inference of these parameters by Beyeler et al. (2019b) based on four patients. For the axons shape descriptors, we relied on the range of values determined by Jansonius et al. (2009) on 65 subjects. For the range of the remaining parameters (implant position and orientation), no data were available about postoperative implant movement at the time of the experiment, so we choose a realistic value range. The bounds of the search space are summarized in table 2.1.

	ρ (μm)	λ (μ m)	<i>x</i> (µm)	<i>y</i> (µm)	θ	β_{-}	eta_+	Amplitude
Lower bound	137	358	-500	-500	-15°	0.1	-2.5	0
Upper bound	415	1510	500	500	15°	1.3	-1.3	300

Table 2.1: Search space. ρ is the current diffusion coefficient, λ is the axon sensitivity to current. ρ and λ ranges are taken from Beyeler et al. (2019b). x and y are the implant coordinates with respect to the fovea. θ is the implant orientation with respect to the horizontal axis. β_{-} and β_{+} are axons shape descriptors, their values are taken from Jansonius et al. (2009). The amplitude parameter scales the percept brightness.

Kernel choice and hyperparameters selection

The hyperparameters and the kernel of the surrogate model can have a dramatic effect on performance (Snoek et al., 2012). They are usually adapted over the course of the optimization, using Bayesian model selection (Rasmussen and Williams, 2006), but this comes with several drawbacks. First, estimating the hyperparameters at each iteration slows the acquisition process down, but one of the criteria we had for our optimization strategy is that it should be fast enough so that the time between two queries would be much shorter than the query itself. Second, with a small amount of data, maximum likelihood estimation usually leads to overfitting, severely affecting the optimization success.

To overcome these difficulties, we finally adopted a transfer learning strategy, where one participant was asked to run 400 pairwise comparisons between random encoders, with two different perceptual models chosen at random. The data obtained with the first perceptual model can be seen as a training set, and the other as a test set. We then trained GP models on the training set, either with the Squared Exponential with Automatic Relevance Determination (SE-ARD), Matérn 3/2, or Matérn 5/2 kernel (see appendix C for the definitions of these kernels). We inferred the hyperparameters using type-II maximum-likelihood. We then tested the GP models on the test set and selected the best performing kernel by measuring each model's Brier score (a widely used measure of probabilistic predictions accuracy). This kernel was later used In the experiment, with the hyperparameters learned on the training set. Interestingly, the best kernel on the training set was SE-ARD, but on the test set, it was the Matérn 5/2 kernel.



Figure 2.2: Principle of hyperparameters transfer. Two datasets of 400 comparisons between encoders are built using two different perceptual models. GP models are trained on the training dataset with three different kernels, and for each kernel, the hyperparameters are inferred using type-II maximum likelihood. After this model selection step, GP models are learned using the same hyperparameters on the test data set. The kernel for which the Brier score is the highest on this second data set is used in the experiment, with the same hyperparameters.

This probably comes from the fact that the SE-ARD kernel has more parameters (9 parameters, but only 2 for the Matérn kernels), which makes the GP model a better description of the data. This strategy is illustrated on figure 2.2. This can be seen as a way of transferring knowledge from previous patients to new ones.

2.3.4 Preference and visual acuity measurements

To determine whether the encoder found after running our optimization procedure was indeed preferred to other encoders and led to improved visual function, we measured preference and visual acuity at the end of the experiment.

Preference To measure preference between two encoders, we jointly showed on the monitor the percepts predicted for two different encoders for a given stimulus and asked the subject to indicate the one he found the least distorted (figure 2.1). To get an accurate measure of the subject's preference, we repeated the comparison for 13 random letters. The name of the letter was indicated above so that the subject knew what the reference stimulus was. The letters either came from the same set of 13 letters that were used during the optimization process (denoted as "optimization set") or from another set (denoted as "transfer set").

Visual acuity To measure visual acuity, we used two distinct tasks: the tumbling E, and the Snellen chart task (including the letters C, D, E, F, H, K, N, P, R, U, V, and Z).

In the tumbling E task, we presented the letter E at various sizes at four different orientations and asked the subject to report the perceived orientation. In the Snellen chart task, we presented letters of various sizes and asked the subject to identify the letters.

We fitted the participant's responses in visual acuity measurements using a psychometric

function:

$$P(\text{Correct response}) = \gamma + (1 - \gamma)\Phi(as + b)$$

Where *s* is the stimulus finest detail size (1/5 of the size of the letter) in log of the visual angle in minutes of arc, γ is the chance level, *a* and *b* are constants. The visual acuity is the inflection point of the curve, which also corresponds to the smallest size for which *P*(Correct response) = $\frac{\gamma+1}{2}$.

To efficiently measure these thresholds, we used the adaptive psychophysics procedure QUEST+ (Watson, 2017; Jones, 2018). The size of the visual field in the prosthetic vision simulations was $21^{\circ} \times 29^{\circ}$, so the maximum size of the letter was 21° .

Real-life patients use scanning when performing visual acuity tests, which likely leads to overestimated visual acuity scores (Ayton et al., 2014). Moreover, patients have extensive experience of prosthetic vision, so care should be taken when comparing the visual acuity measured in our experiments to the performance of real-life patients.

2.3.5 Bayesian analysis

To analyze the data from preference comparisons between encoders after the experiment, we used the following Bayesian analysis method adapted from Kass and Raftery (1995) (see also chapter 5 of Robert (2006)): For two algorithms A_1 and A_2 used to derive encoders, we have the two hypotheses H_1 (resp. H_2): encoders found using A_1 (resp. A_2) are preferred to encoders found using A_2 (resp. A_1). For a participant *i*, the responses in a series of preference comparisons between two encoders follows a binomial distribution $y_i = \mathcal{B}(n_i, p_i)$, where n_i is the number of comparisons and y_i is the number of comparisons for which the participant reported that A_2 is better than A_1 . We put the conjugate beta prior $p_i \sim Beta\left(\frac{\xi}{\omega}, \frac{1-\xi}{\omega}\right)$, where $\xi \in [0, 1]$ and $\omega \in \mathbb{R}^*_+$. Note that $\xi = \mathbb{E}(p_i)$. We assume that ω is constant, and ξ follows a uniform distribution. Under H_1 , the model is $\xi \sim \mathcal{U}([0, 0.5])$, whereas under H_2 , $\xi \sim \mathcal{U}([0.5, 1])$.

The Bayes factor B_{21} is the ratio between the likelihoods of the two models: $B_{21} = \frac{p(\mathscr{D}|H_2)}{p(\mathscr{D}|H_1)}$. It gives an estimate of how strongly the data support H_2 over H_1 . The details of the computation of $p(\mathscr{D}|H_i)$ are described in appendix A.

For each comparison between algorithms, ω was estimated empirically, by noting that: $\mathbb{V}(p_i) = \frac{\omega\xi(1-\xi)}{1+\omega}$. So we computed the empirical mean of the p_i , $\hat{\xi} = \langle p_i \rangle$ and the empirical variance $\hat{V} = \langle p_i - \hat{\xi} \rangle$, and solved for $\hat{V} = \frac{\omega\xi(1-\xi)}{1+\omega}$. We then evaluated the integral numerically. An advantage of this analysis method is that it models the noise in the preference measurement data.

To compare visual acuity in the E and Snellen tasks between different encoders, we assumed that the difference in visual acuity Δ VA measured between encoders obtained with two methods *A* and *B* follows $\mathcal{N}(\mu_{AB}, \sigma_{AB}^2)$. The two competing hypotheses are H₁: $\mu_{AB} > 0$, and H₂: $\mu_{AB} < 0$. We infer σ_{AB}^2 by measuring the empirical variance, and we put a uniform prior $\mu_{AB} \sim \mathcal{U}(a_i, b_i)$, where $a_1 = b_2 = 0$, $b_1 = v$, and $a_2 = -v$, where *v* corresponds approximately to the maximum of visual acuity that we can measure.

We have:

$$p(\mathscr{D}|H_i) = \int \prod_i \mathcal{N}(\Delta \mathrm{VA}_i | \mu_{AB}, \sigma_{AB}) \mathcal{U}(\mu_{AB} | a_i, b_i) d\mu_{AB}$$
(2.9)

Based on 2.9, the Bayes factor can be computed via Monte-Carlo integration.

One advantage of using this statistical analysis method is that the log of the Bayes factor is symmetric. A $K = \log_{10}$ (Bayes factor) above 0.5, 1 or 2 corresponds to substantial, strong and decisive evidence in favor of H₂ respectively. Conversely, a $K = \log_{10}$ (Bayes factor) below -0.5, -1 or -2 corresponds to substantial, strong and decisive evidence in favor of H₁ respectively.

2.4 Results

We simulated that each of our participants had been fitted with an Argus II prosthetic device, but with different parameters describing their perceptual model. To do so, a set of parameters ϕ^* was randomly assigned to each participant, defining the mapping from electrical pulses to percepts. These parameters were sampled uniformly in a range of realistic values corresponding to the search space of the optimization (see table 2.1). The only parameters that were kept constant across sampled perceptual models were the axon shape descriptors, which define the trajectories of axons in the retina. These two parameters were set at extreme values as they were later modified to assess the effects of a wrong parameterization (see part 2.4.5). Apart from amplitudes, the pulse parameters were kept at their standard values in the *pulse2percept* model (Beyeler et al., 2017a). The simulated visual field size was 21×29 degrees of visual angle. For the fraction of the participants' visual field to approximately correspond to this, we asked them to sit at the corresponding viewing distance.

2.4.1 Model-based encoding of visual inputs

The standard approach to visual signal encoding in visual prosthetics is to downsample the input image and to linearly convert resulting brightness to currents amplitudes (Bloch and da Cruz, 2019). We first wanted to assess to which extent does the model-based parameterization of the encoder improve the perceptual experience of participants compared to this naive approach. In practice, we computed each electrode's receptive field by tesselating the input visual space so that each pixel drove only one electrode. This electrode was the one closest to the pixel (this is known as Voronoi tesselation). The linear scaling was adjusted so that the predicted percept would not saturate.

We then asked a subset of 14 participants to compare the experienced quality of the naive encoder to a model-based encoder. If we knew the correct parameters of each patient's perceptual model, then we could minimize distortion exactly, but here, we wanted to see if the mere model-based parameterization of the encoder would improve vision. We thus selected random parameters. Interestingly, participants showed strong preference for the parameterized encoder, even though its parameters were chosen at random (figure 2.3, log₁₀ of the Bayes factor: K = 2.4, decisive evidence). Similarly, the visual acuity was significantly improved when using the model-based encoder compared to the naive encoder, in both the tumbling E task (log₁₀ of the Bayes factor: K = -2.2, decisive evidence) and the Snellen task (figure 2.3, K = -0.85, substantial evidence). These results suggest that model-based parameterization of visual prosthetics is effective in improving the prosthetic users' visual experience. However, when comparing predicted percepts corresponding to an encoder with random parameters to those corresponding to the "optimal" encoder (the one based on the ground truth perceptual


Figure 2.3: A. Percept predicted by the pulse2percept model with the naive encoder, for an input stimulus corresponding to the letter A. B. Percept predicted by the pulse2percept model with the control encoder, which corresponds to random parameters ϕ , for the same stimulus as A. C. Preference between the control and the naive encoder. In most cases the participants prefer the control encoder (K = 2.4, N = 14, decisive evidence). D. VA measured in the tumbling E and Snellen task with the control or the naive encoder. The control encoder corresponds to better visual acuity in both the tumbling E (K = -2.2, N = 14, decisive evidence) and the Snellen task (K = -0.85, N = 14, substantial evidence). The asterisks represent the centroids.

model parameters), it appears that the encoder parameters can drastically affect the perceptual quality (figure 2.4.A).

2.4.2 Preference-based optimization of retina implant encoder

Of course, ideally, we would like to select parameters ϕ that maximize some measure of visual function in a subject-specific manner, rather than picking them at random. As described earlier, we did this for each subject using preference-based optimization. In practice, we followed the preference measurement procedure described above, but for a given pair of encoders, we only presented one comparison (figure 2.1). The stimulus was chosen at random among a set of 13 letters (the optimization set).

Then, the closed-loop optimization procedure was run during 60 iterations. The experiment was repeated either using the adaptive parameters selection procedure (that is, using the maximally uncertain challenge acquisition rule) or using a non-adaptive procedure where duels were selected at random.

Figure 2.4.B shows the percepts presented during a typical experiment and the ones corresponding to the inferred best encoders. As data accumulate, the duels correspond to less distorted percepts, and the inferred percept also improves. Figure 2.4.A shows the percepts corresponding either to the encoder computed with the ground truth subject's parameters, the encoder optimized using the adaptive procedure, the encoder optimized using a non-adaptive acquisition (where at each iteration, the pair of encoders was selected at random), and an encoder with random parameters. In this example, the percept looks less distorted with the encoder derived from the optimization sequence that was based on adaptive optimization compared to non-adaptive optimization. To get a quantitative measure of the perceptual quality with different encoders, we assessed preference between encoders on the set of stimuli used during the optimization process.



Figure 2.4: A. Predicted percepts with the encoder corresponding to the true parameters, the encoder optimized using the adaptive optimization procedure, the encoder optimized using the non-adaptive method, and the control encoder corresponding to random parameters for the same participant. The percepts tend to be discontinuous at the limit between the upper and lower hemifields, this is due to the geometry of the axons. B. The two left columns correspond to duels presented during the optimization sequence (the champion is on the left, and the challenger is on the right, but in the experiment, they were positioned randomly). The far-right column corresponds to the percepts predicted with the inferred best encoder at the corresponding iteration, with the letter A as input stimulus. C. Preference ratios measured on the set of stimuli used during the optimization, between the encoder optimized using adaptive preference optimization procedure, compared to encoders with random parameters, the encoder optimized using the non-adaptive procedure, or the ground truth encoder. The encoder found using the adaptive procedure is preferred to the encoder with random parameters (K = 16, decisive evidence), and to the encoder derived from the optimization based on non-adaptive acquisition (K = 0.8, substantial evidence). However, the encoder based on the true perceptual model parameters (ground truth) is significantly preferred on average (K = -1.9, strong evidence). **D.** Preference ratios on the set of stimuli that were not used during the optimization process. The conclusions are identical, except that on this set of stimuli there is barely a noticeable difference between encoders derived from the adaptive or non-adaptive *procedure* (K = 0.45).

We asked a subset of 20 subjects to compare the encoder optimized using the adaptive preferential procedure to encoders corresponding to random parameters (figure 2.4.C). All participants preferred the optimized encoder, demonstrating that our algorithm can efficiently improve the encoder's quality.

To evaluate whether the adaptive procedure led to better encoders compared to the nonadaptive procedure, where at each iteration pairs of encoders were selected at random, all 24 participants were asked to compare the two types of encoders. The results show that, on average, the adaptive procedure leads to preferred encoders (K = 0.80, substantial evidence). An adaptive optimization strategy is thus beneficial. However, the encoder computed using the ground truth perceptual model parameters was preferred to the optimized ones (K = -1.9, strong evidence), meaning that our optimization procedure did not converge to the best possible encoders in 60 iterations.

The fact that encoders optimized to minimize the mean-squared error between stimuli and percepts are preferred to encoders optimized using adaptive preferential optimization suggests that MSE could be a good metric for perceptual quality. In order to assess whether this is true in general,we computed, for each encoder, the MSE between predicted percepts and stimuli for the 26 letters (figure 2.5). Interestingly, this distance is higher for the optimized encoders than for the encoders with random parameters (figure 2.5.A, K = 5.4, decisive evidence), suggesting that optimization would have led to different results if it had been based on MSE. Because we simplified the MSE minimization problem in equation 2.2 by removing the non-linearity, the encoder parameterized by ϕ^* is not the true minimizer of the MSE. However, as seen in figure 2.5.B, in most cases it leads to smaller MSE with respect to the stimuli than the optimized encoders (K = 10, decisive evidence). More generally, the difference in MSE between encoders does not depend on the preference between them (figure 2.5.C, F-test, p = 0.86), suggesting that this metric does not reflect perceptual quality. The fact that the value that a subject attributes to a specific encoder cannot be predicted by the corresponding MSE implies that subjects feedback is required to optimize the encoder.

2.4.3 Transfer to other tasks

Preferential Bayesian optimization is efficient at improving perceived images quality in simulated prosthetic vision. A fundamental requirement for our strategy to be useful to patients is its ability to generalize improvement to other tasks than the one for which the encoder was optimized. To assess whether perceptual improvements are transferred to other tasks, we first measured preference between encoders on another set of stimuli, comprised of 13 letters that were not presented during the optimization phase (figure 2.4.D). Again, optimized encoders were always preferred to encoders with random parameters. Although these stimuli still correspond to high contrast letters, this shows that the optimization generalizes beyond the optimization set.

Then, we asked whether optimization led to improved performance in two tasks commonly used to measure visual acuity. The 'tumbling E' task (figure 2.6.A), where the goal is to detect the orientation of a square E among four possibilities, and the 'Snellen chart' task (figure 2.6.B), where the subject has to recognize a letter.

We measured a significant improvement in visual acuity compared to encoders with ran-



Figure 2.5: We computed, for each encoder, the MSE between predicted percepts and stimuli for the 26 letters. **A.** The MSE is higher on average for the optimized encoder than for the encoders with random parameters (K = 5.4, decisive evidence). **B.** The MSE is higher on average for the optimized encoder than for the encoders with the true model parameters (K = 10, decisive evidence). The asterisks represent the centroids. **C.** The difference in MSE between two encoders is not correlated with the preference between them (F-test, p = 0.86).

dom parameters (figure 2.6.C), in both the E task (K = 4.9, decisive evidence) and the Snellen task (K = 3.0, decisive evidence). This demonstrates that improvements transfer to tasks very different from the ones used in the optimization process.

The performance was slightly improved in the tumbling E task when using encoders derived from the adaptive optimization procedure, where candidate parameters were carefully chosen to maximize information gained from subjects, versus the non-adaptive one, where parameters were selected randomly at each iteration (figure 2.6.D, K = 0.54, substantial evidence). However, in the Snellen task, we could not detect significant differences. Similarly, when measuring preference on the new set of stimuli, there was no marked preference for encoders derived from the adaptive procedure compared to the non-adaptive one.

When comparing VA between the encoder optimized using the adaptive procedure and the encoder corresponding to the true subject's parameters, we observed that the ground truth encoder outperformed the optimized encoder (K = -4.8 and K = -2.2, decisive evidence), which is consistent with the results on preferences.

2.4.4 Comparison with performance-based optimization

We were interested to compare our preference-based learning approach to an alternative approach, based on optimizing subjects' performance in a given task. In order to compare the two strategies and determine their relative advantages, we ran a second set of experiments in which we tried to maximize subjects' performance in a tumbling E task for constant letter size (the letter covered 7.4° of visual field, this size was chosen in pilot studies to get a success rate of about 0.5 with encoders with random parameters). The optimization framework is the same, as both preference and performance in forced-choice tasks can be modeled using classification models. As an acquisition function, we used Thompson Sampling, a widely used heuristics in Bayesian optimization (Thompson, 1933). Thompson sampling consists in selecting as the next



Figure 2.6: Visual acuity measurements following preference-based optimization. A. The tumbling E task consists in identifying the orientation of a letter E among four possibilities. **B.** The Snellen chart task consists in recognizing a letter from the Snellen letters chart. We measured visual acuity with a specific encoder in each of these tasks by varying the letter size. The visual stimulus was encoded into a current amplitudes pattern and the evoked percept was simulated using the pulse2percept model. The size of the letters was adaptively varied for efficient visual acuity estimation using the adaptive psychophysics procedure QUEST+ (Watson, 2017; Jones, 2018). C. Encoders found using the adaptive preferential optimization procedure (adaptive pref.) lead to better visual acuity compared to encoders with random parameters (random ϕ) in both the tumbling E (K = 4.9, decisive evidence) and the Snellen chart task (K = 3, decisive evidence). The asterisks represent the centroids. **D.** However, there is only a slight improvement in the tumbling E task with these encoders compared to encoders derived from the non-adaptive preferential optimization (non-adaptive pref., K = 0.54, substantial evidence), and no significant difference in the Snellen chart task (K = -0.13). E. The encoders found using the adaptive preferential optimization procedure do not match the performance of the encoders based on the true perceptual model (ground truth) in both the tumbling E task (K = -4.8, decisive evidence) and the Snellen chart task (K = -2.2, decisive evidence).



Figure 2.7: A. Comparison between visual acuity measured with the encoders found with the performance-based optimization procedure and encoders with random parameters (random ϕ). Their is no significant improvement after the optimization sequence, neither in the tumbling E task (K = 1.3×10^{-1}), nor the Snellen task (K = -3.0×10^{-1}). The asterisks represent the centroids. **B.** Preference based-encoders obtained with the adaptive preferential optimization procedure (adaptive pref.) lead to better visual acuity in both the tumbling E (K = -3.0, decisive evidence) and the Snellen task (K = -3.1, decisive evidence). **C.** Performance-based encoders are not significantly preferred to encoders with random parameters (K = 0.10). Preference-based encoders are preferred to performance based-encoders on the set of stimuli that was not used to optimize the encoders (K = -7.38, decisive evidence).

query a sample *x* from $p(x^*|\mathscr{D})$.

Performance-based optimization did not lead to improved visual acuity compared to randomly sampled encoders (figure 2.7.A). Preference-based optimization procedure lead to a better visual acuity (figure 2.7.B) in both tumbling E (K = -3, decisive evidence) and Snellen chart tasks (K = -3.1, decisive evidence) compared to performance-based optimization. Finally, when measuring preference between the two encoders (figure 2.7.C), preference-based encoders are strongly preferred to performance-based ones (K = -7.8, decisive evidence), whereas performance-based encoders are not preferred to encoders with random parameters.

The preference-based strategy is thus much more efficient than the performance-based approach we tested here. Moreover, the preference-based approach requires much less calibration, making it easier to apply. Indeed, the performance-based optimization experiments required tuning the difficulty of the task based on the expected performance of subjects. Finally, the participants frequently reported they found the preference judgment task much more engaging, which is of practical importance for patients who have to engage in prolonged rehabilitation protocols.

2.4.5 Robustness to mis-parameterization of the encoder

When applying our strategy to real-life patients, there will likely be a mismatch between the perceptual model based on which the encoder is parameterized and the patient's true perceptual model. To estimate the impact of such misspecification, we ran the same experiments



Figure 2.8: A. Axon map in the true patient model. The shaded area corresponds to the location of the implant. **B.** Axon map in the misspecified model. Note how the axon trajectories are affected, in particular in the region below the implant. **C.** Example of percept predicted with the encoder derived from the true perceptual model (correct model) or with the encoder derived from the misspecified perceptual model, where all parameters were set at their true value, except the two axon shape descriptors (misspecified). Note that the percept quality is severely affected by the model misspecification. **D.** Effect of axon shape descriptors misspecification on visual acuity. With all parameters at their true value, except the two axon shape descriptors, visual acuity is significantly reduced in both the Snellen chart task (K = -13, decisive evidence) and the tumbling E task (K = -8.6, decisive evidence) compared to the encoder based on the ground truth parameters. The asterisks represent the centroids.

keeping the parameters corresponding to axons shape descriptors fixed. These values were chosen to be as different as possible compare to the values used to define the patient perceptual model (figure 2.8.A and 2.8.B). We choose to misspecify these two parameters are they are known to vary from one subject to another (Jansonius et al., 2009), were not inferred in the original Beyeler et al. (2019b) study, and are a priori impossible to measure in an already implanted patient. We confirmed that the misspecification of the encoder led to severely degraded percepts (figure 2.8.C), linked to poor performance in all vision tests compared to the encoder derived from the true perceptual model (figure 2.8.D, K = -8.6 and K = -13 for the E and Snellen task respectively, decisive evidence).

Adaptive preference-based optimization with this misspecified model led to optimized encoders that were preferred compared to encoders with random parameters and misspecified axons shape descriptors (K = 1.5, strong evidence), as well as to encoders with random parameters (K = 1.6, strong evidence) (figure 2.9.B). We conclude that misspecification does not impede the preference-based optimization process.

However, these optimized encoders were not preferred to the ones derived from the model in which all but the axon descriptors were at their true value (figure 2.9.B). This suggests that the optimization of the free parameters is not able to fully compensate for the misspecification of axon shape descriptors.

Performance in the tumbling E and Snellen tasks was also better after optimization, despite misspecification, compared to encoders for which parameters were selected at random (figure 2.9.C) in both the tumbling E task (K = 1.8, strong evidence) and the Snellen chart task (K = 0.55, substantial evidence). This demonstrates that the transfer of improvement to other tasks



Figure 2.9: Robustness of the optimization process with a misspecified model. A. Example of percepts predicted for the stimulus in figure 2.8.C, either with the encoder based on the misspecified model (misspecified) or with the encoder optimized using preferential Bayesian optimization with a misspecified model (opt. miss.). In this example, the optimization procedure found an encoder that partly compensated for the model misspecification. **B.** Encoders optimized based on a misspecified model were preferred to encoders where parameters were selected at random, except the axon shape descriptors that were misspecified (random ϕ miss., K = 1.5, strong evidence), and to encoders with random parameters (random ϕ , K = 1.6, strong evidence). However, these encoders were not preferred on average to encoders found by inverting the misspecified model (misspecified, K = 0.25). This shows that despite the robustness of the optimization procedure, the limited encoder flexibility did not compensate for the misspecification. **C.** Visual acuity was better with encoders optimized based on a misspecified model (opt. miss.) compared to random encoders in both the tumbling E task (K = 1.8, strong evidence) and the Snellen chart task (K = 0.55, substantial evidence). The asterisks represent the centroids.

is not impeded by the model misspecification.

Overall, the fact that neither the optimization process nor the transfer of performance improvement to new visual tasks are significantly affected by the perceptual model misspecification demonstrates the robustness of our method.

2.5 Discussion

2.5.1 Summary of contributions

In this paper, we provided a proof of principle for a new strategy to shape the phosphenes experienced by prosthetic vision users. Our method is readily applicable to real-life patients, and would solely require implementing our algorithm in the visual processing unit.

This strategy is based on two pillars: first, inversion of a model of prosthetic vision to parameterize an encoder. Second, optimization of the encoder parameters using patients' subjective reports. As an optimization algorithm, we used preferential Bayesian optimization, an efficient method that builds a model of the subject's preference to adaptively select new configurations to test. We assessed our strategy on healthy subjects using a state-of-the-art prosthetic vision simulator (Beyeler et al., 2017a).

We demonstrated the efficiency of global phosphene shaping in improving the perceptual experience in simulations of prosthetic vision. Moreover, we showed that preferential Bayesian

optimization of the encoder led to significant improvement in perceptual experience compared to randomly selected encoder parameters. Importantly, this improvement transferred to tasks that were distinct from the ones used in the optimization procedure, leading to higher visual acuity. Notably, the adaptive optimization procedure led to better results than the non-adaptive one. We also observed that, in general, the preference between two encoders was not related to how close the resulting percepts were from the original stimuli (as measured using mean-squared error). This suggests that an open-loop optimization based on minimizing the MSE, or a closed-loop optimization targeting a specific RGCs activity pattern minimizing MSE between predicted percept and stimuli (Shah et al., 2019) may not exploit the full capacity of the prosthetic device's sight restoration potential.

Our framework also allows leveraging prior information from other patients. Indeed, Bayesian optimization allows taking into account prior knowledge about the objective function. So information from previous patients can be incorporated to guide the optimization process. In our experiments, the search space was constrained by prior knowledge of the distribution of the perceptual model parameters among patients (Perrone et al., 2019), and the hyperparameters of the Gaussian process kernel were learned in one participant and kept fixed for the others.

2.5.2 Related work

The idea of inverting a model of prosthetic stimulation so as to shape the electric field has been discussed in Spencer et al. (2018). The goal was to control the retina ganglion cells activation pattern assuming a linear-nonlinear model estimated from recordings made with the implant. This strategy has also been applied to transcranial direct current stimulation (tDCS) in order to get intense and focal stimulation (Dmochowski et al., 2011). Chen et al. (2020) followed a very similar approach in retinal prostheses, but instead of optimizing the stimulation based on an estimate of the tissue response, the optimization was based on the electrodes induced electric fields. The proposed stimulation strategies effectively invert the model to find the pattern of electrical stimulation that optimally matches a target pattern of activity. Our approach differs as we focused on directly shaping the phosphenes based on a model of prosthetic vision, the advantage being that this does not require retinal activity recording (which is impossible with current devices).

A closed-loop optimization of retina implant encoders was proposed in a series of papers by Eckmiller et al. (1999); Becker et al. (1999); Eckmiller et al. (2005, 2007). Subjects had to compare the perception elicited by a given stimulation to an expected percept, and new parameters were iteratively proposed according to genetic or reinforcement learning algorithms. However, this strategy was not demonstrated on a prosthetic vision model validated with patients' data. Moreover, the authors did not show that the optimization procedure led to improvement in sight restoration, and they do not mention whether the stimulation parameters optimized for one input pattern also improved perception of other shapes. This approach was based on a perceptual comparison between an input pattern and a perceived pattern, but real-life patients do not have access to the input pattern. Our approach, on the contrary, is based on a comparison between percepts, and the reference to the cue is symbolic (we displayed the name of the letter corresponding to the stimulus, but in real-life patients, the letter could be indicated orally). Finally, the optimization procedure we used led to significant improvements in the perceptual experience in a much shorter time (only a few minutes, whereas the procedure by Eckmiller et al. (2007) took 60 to 120 minutes).

2.5.3 Performance-based vs preference-based optimization

Our rationale for using preference-based optimization came from the idea that it allows to take into account and balance aspects of perception difficult to grasp with performance measurement tasks. However, we could not exclude that the subjectivity of feedback in preferencebased optimization may lead to a perceptual experience that the subject finds more qualitative, but which does not improve general visual function. In reality, our results support the hypothesis that qualitative experience improvement also leads to improvement in objective performance in acuity-related tasks.

By contrast, a fundamental limitation of performance-based optimization is that improvement in the task, or set of tasks, may not be related to vision improvement in general. Indeed, the optimization process may emphasize artefactual perceptual cues that are not related to vision. This problem is related to the more general vision testing problem in the field of vision restoration and has been extensively discussed by Peli (2020). As a consequence, we expected preference-based optimization to generalize better than a single performance-based optimization. We could not test this hypothesis as the performance-based optimization we assessed did not improve subjects' visual function. Multiple reasons may explain that the performancebased optimization failed. First, performance measurements in binary tasks may be less informative than pairwise comparisons. Second, the task may be either too difficult or too easy so that responses are not informative. If the encoder improves over the course of the optimization process, without varying the difficulty of the task, the optimization algorithm will not differentiate between encoders that are equally good for that level of difficulty. At the time we did the experiments, we could not find existing optimization algorithms addressing the problem of tuning the task difficulty (we address this problem in chapter 4). Preference-based optimization has a similar drawback. For a given set of stimuli, the subject may find all the encoders that are proposed to him equally good or equally bad. Again, this highlights the importance of choosing appropriate stimuli. In our experiment, however, we did not have to hand-tune the set of stimuli. This suggests that designing a preference-based optimization protocol is much easier than selecting forced-choice visual tasks.

Of course, the two approaches are not exclusive and could potentially be combined to balance the advantages of using subjective and objective criteria.

2.5.4 Robustness of the adaptive preferential optimization procedure

Notably, our demonstration avoids a common pitfall of prosthetic stimulation optimization research, in which the open-loop optimization of the electrical stimulation is based on a known decoder (see e.g. Nirenberg and Pandarinath (2012)). Here, we avoid this circular reasoning by showing that our closed-loop strategy is robust to encoder-decoder mismatch. Moreover, these approaches are based on the hypothesis that we have access to a metric relating RGCs activity to visual function (Shah et al., 2017), typically assuming that the most natural the neural response

is, the better. With our closed-loop strategy, we do not explicitly make use of a metric and directly ask the subjects to compare settings. Our results support this approach, as we showed that preference between encoders is not related to their difference in mean-squared distance between stimuli and percepts.

Our preferential Bayesian optimization strategy harnesses the power of prosthetic vision models for dimensionality reduction and statistical learning for robustness and flexibility. As a consequence, the optimization process is not impeded by the misspecification. This suggests that our method could be applied in real-life patients even if there is a severe mismatch between the model used to parameterize the encoder and the true perceptual model. However, our results show that optimizing the parameters does not compensate for the misspecification, so that the device does not reach its full vision restoration potential. This is due to the fact that by parameterizing the encoder, we restrict the search of the optimum to a subspace in the space of all possible encoders. This makes the optimization easier but limits flexibility.

2.5.5 Limits and future developments

Likely, more accurate prosthetic vision models will also include more parameters. While this could improve the extent to which visual function can be improved thanks to increased encoder flexibility, this would also make the optimization problem more difficult. In general, for a given amount of data, the encoder dimension determines a tradeoff between the encoder ability to flexibly adapt to the patient's preference, and the difficulty to learn this preference as dimension increases. In particular, highly flexible encoders would require enhancing the Bayesian optimization algorithm with features specifically designed for high-dimensional optimization (see e.g. Gardner et al. (2017); Wang et al. (2017b); Mutný and Krause (2018); Rolland et al. (2018); Li et al. (2017); Zhang et al. (2019)).

To limit the total duration of the experiment for individual participants, we restricted optimization sequences to 60 iterations. At each iteration, computing the new candidate encoders required to run the *pulse2percept* model, which took a few seconds, so that an optimization sequence lasted about 8 minutes. Our results show that despite considerable improvement in encoder quality, this was not sufficient to converge to the preferred encoder. In application to real-life patients, an optimization sequence could probably last a few hours, allowing the encoder to improve even further.

We did not consider the temporal dynamics of stimulation and percepts. Electrical stimulation induces a desensitizing phenomenon in RGCs, whereby phosphenes gradually fade over time (Horsager et al., 2009; Freeman et al., 2011; Pérez Fornos et al., 2012). Our approach could generalize to include the temporal dimension in the encoder and in the tasks. For example, we could ask the subjects to compare their perceptual experience between encoders for static stimuli exactly as we did, but the subjects would also have to consider the temporal dynamics of their percept (for example, a subject would likely prefer a percept that does not fade to one that does). Our approach could also be used to tune stimulation parameters other than currents amplitudes, such as pulses shapes and frequency, for more precise tuning.

Our model-based parameterization of the retinal prosthetic encoder corresponds to minimizing the expected distance between the stimulus and the percept (see equation 2.1). In practice, however, the distortion minimization problem is constrained due to the charge injection limit of metallic electrodes. The Argus II Surgeon Manual (Second Sight Medical Products: Argus II Retinal Prosthesis System Surgeon Manual (2013)) specifies that the maximum current per electrode is 1.0 mA. However, the *pulse2percept* model predicts that the percept brightness scales approximately linearly with the pulse frequency (Nanduri et al., 2012). That way, it is possible to reduce pulse amplitudes while increasing the frequency to stay below safety limits. Another, more general approach to enforce such constraints on the encoder would be to use a quadratic optimization algorithm with linear constraints.

In general, for more complex encoders including varying pulse shapes, frequencies, etc, which involve non-linear effects, more advanced optimization techniques would be required to invert the perceptual model. Note, however, that this is fully compatible with our approach, the only difference being the time taken to compute a given encoder.

In our experiments, the preference-learning task was general enough so that the optimized encoder was appropriate in a range of different tasks. A possible implementation would specifically adapt the encoder to different activities or environments. One could also consider adaptively varying the stimuli that are presented during the optimization sequence, to incorporate more complex stimuli with finer detail. Hopefully, this would allow the user to make refined comparisons, and enhance the transfer of improvement to other tasks.

A significant advantage of our approach is that subjects could optimize their retina implant encoder on their own during daily life, without expert assistance. This could be particularly useful as the retina continues to degenerate over time, the implant/retina interface evolves, and the subject learns to make use of prosthetic vision (Beyeler et al., 2017b). The long-term influence of artificial stimulation of the retina on network architecture is still unknown and is a remaining concern, but our system could circumvent those changes by adapting the way the electrode array transmits visual stimuli.

2.5.6 Code availability

The Matlab code to perform the experiment and run the analyses is available at https://github.com/TristanFauvel/Retinal_Prosthetic_Optimization_Code.git. The implementation of the QUEST+ adaptive psychophysics procedure is from Jones (2018). To display stimuli and record responses, we used the Matlab toolbox Psychoolbox (Brainard, 1997; Pelli, 1997; Kleiner et al., 2007)

2.6 Chapter supplementary

2.6.1 Instructions to the participants

The exact instructions given to the participants were the following (eventually translated to French for native French speakers):

- Preference comparisons: 'The image corresponds to a distorted version of a letter (printed above), indicate which one you find the most recognizable or least distorted using the left or right arrows.'
- Tumbling E task: 'The image corresponds to a distorted version of a letter E, which can take 4 different orientations. Indicate the orientation by pressing the corresponding arrow keys.'
- Snellen chart task: 'The image corresponds to a distorted version of a letter. Press the keyboard to identify the letter.'



2.6.2 Supplementary figures

Figure 2.10: Comparisons between percepts predicted by the pulse2percept model for the stimulus in \mathbf{A} , either obtained with the analytical solution of the simplified problem 2.1 (in \mathbf{B}), $\mathbf{M}_{\phi}^{\dagger}$, or the encoder that is a numerical solution of the exact minimization problem \mathbf{W}_{ϕ} (in \mathbf{B} , note that we specifically optimized the amplitudes of the 60 electrodes for this stimulus). These two solutions lead to virtually identical percepts. **D**. Comparison of the distance (RMSE) between the stimulus and the percept obtained using either $\mathbf{M}_{\phi}^{\dagger}$ or \mathbf{W}_{ϕ} for 1000 randomly chosen perceptual models. **C**.Histogram of the RMSE between percepts obtained either with $\mathbf{M}_{\phi}^{\dagger}$ or \mathbf{W}_{ϕ} for 1000 randomly chosen perceptual models.

2.6. Chapter supplementary



Figure 2.11: Comparison of results for participants who participated to the experiment remotely (in blue) or in the lab (in red). A. Visual acuities measured in the tumbling E task. B. Visual acuities measured in the Snellen task. Visual acuities do not differ significantly between groups (two-samples t-test at 0.05 significance threshold) C. Preference comparisons between encoders. There is no contradiction between the two groups.

3

Practical Bayesian optimization with binary and preference data

In the previous chapter, we applied Bayesian optimization to the optimization of retinal prosthetics encoders. Bayesian optimization is a powerful method for optimizing black-box functions, that is, functions whose analytical form is unknown, and for which no derivative information is available (see Brochu et al. (2010b); Shahriari et al. (2016); Frazier (2018) for review). It is especially relevant when evaluation of the function is expensive, either in time, computational resources, price, etc. In particular, Bayesian optimization has been extensively applied to the problem of tuning machine learning model hyperparameters (Turner et al., 2021), and to human-machine interfaces such as exoskeletons (Ding et al., 2018) or deep brain stimulators (Louie et al., 2021). When working with humans, it is often easier to collect binary or preference measurements than absolute rating. For example, when tuning the gait parameters of an ankle-foot prosthesis used by patients with transtibial amputation, it may be easier for the patient to perform pairwise comparisons between settings than to give a grade related to her comfort for a specific setting (Welker et al., 2021). Bayesian optimization has been generalized to this type of scenario, in which evaluation of the function returns binary outputs, such as success/failure or preference judgments. These situations induce several challenges that are not present in standard Bayesian optimization, as the objective function is not directly measured.

In this chapter, we describe the principles of Bayesian optimization with binary outputs and review the various heuristics proposed to select new samples. To the best of my knowledge, there is currently no published exhaustive comparison of existing acquisition rules so that it is difficult for newcomers to determine which one to use. Here, we compare existing acquisition rules on a set of 34 synthetic functions from a widely used set of benchmark functions.

At the core of any Bayesian optimization heuristic is a tradeoff between exploitation (selecting observations at locations were the maximum is likely to lie) and exploration (selecting observations where we are uncertain about the objective value). By revisiting the concept of exploration based on a decomposition of uncertainty into an aleatoric component (relative to the uncertainty on the output that cannot be reduced) and an epistemic component (that reflects our uncertainty about the model), we propose new acquisition functions that outperform all the state-of-the-art heuristics that we tested on the benchmarks, while being fast to compute and easy to implement. We then generalize these acquisition rules to the batch scenario, in which multiple queries are performed simultaneously. Finally, we provide recommendations for efficient implementation.

3.1 Principles of Bayesian optimization

3.1.1 Standard Bayesian optimization

Consider the global optimization problem:

$$\boldsymbol{x}^{\star} = \operatorname*{arg\,max}_{\boldsymbol{x}\in\mathcal{X}} f(\boldsymbol{x}),\tag{3.1}$$

where the objective f is a "black-box" function: that is, we do not know its analytical form, and whether it has special structure like convexity. Moreover, we do not have access to the derivatives of f, and f is typically expensive to evaluate in terms of some limiting resource such as time, money, etc. The input space \mathcal{X} is a bounded set denoted as the search space. We have a fixed budget to evaluate f, that is the maximum number of function queries T is predetermined, and is called the horizon. The regret at time t is defined as $R_t = f(\mathbf{x}^*) - f(\mathbf{x}_t^*)$, where \mathbf{x}_t^* is the maximum among previous observations, that is $\mathbf{x}_t^* = \arg \max f(\mathbf{x})$.

To solve problem 3.1 given a restricted sampling budget, Bayesian optimization proceeds by sequentially querying the function at locations where observations are expected to improve our estimate of the value of x^* . The selection of new values of x to evaluate f follows an acquisition rule, designed to leverage previous data to perform informative observations (Jones et al., 1998).

Concretely, at each iteration, the choice of the next query is made in two steps (see algorithm 1). First, a Bayesian model of the objective function, called a surrogate model or oracle, is learned using previous observations. This surrogate describes the probabilistic belief about the objective and can be used to incorporate prior knowledge. Crucially, the Bayesian surrogate model encodes uncertainty about the underlying objective. Based on this model, a new location to evaluate the objective function is chosen according to a predefined heuristic implemented in an acquisition rule. This rule is designed to efficiently solve the maximization problem. It implements the tradeoff inherent to any global optimization task between spending the evaluation budget in regions of the input space where the surrogate model is highly uncertain (exploration), and regions of the space where the maximum is likely to lie (exploitation). Acquisition rules are usually myopic, i.e., they only consider one-step-ahead choices, without considering the fact that subsequent observations will be performed (Wu and Frazier, 2019).

Due to its generality and performance, Bayesian optimization has become very popular in applications ranging from the design of engineering systems (Forrester et al., 2008) to the tuning of machine learning model hyperparameters (Snoek et al., 2012). Bayesian optimization has been generalized to numerous frameworks, to include optimization with noisy outputs (such as i.i.d. Gaussian noise corrupting the observations, Letham et al. (2019)) or inputs (where there is some uncertainty about the points where the function is evaluated, Fröhlich et al. (2020)), batch optimization (where several points at selected and evaluated simultaneously at each iteration, see e.g. Wu and Frazier (2016)), multi-objective optimization (where several objectives are jointly maximized Emmerich and Deutz (2018)), multi-fidelity optimization (where a trade-off exists between querying the system with high-fidelity but at high cost or with low-fidelity at low-cost, Huang et al. (2006)) and optimization with binary outputs such as preference-based Algorithm 1: Basic pseudo-code for Bayesian optimization

Place a Gaussian process prior on the objective function: $f \sim \mathcal{GP}(\mu(\cdot), K(\cdot, \cdot))$; Observe f at t_0 points according to an initial space-filling experimental design ; Set $t = t_0$; while $t \leq T$ do Update the posterior probability distribution on f using all available data; Let x_t be a maximizer of the acquisition function over \mathcal{X} ; Observe $y_t = f(x_t)$; Augment data $\mathscr{D}_t = \mathscr{D}_{t-1} \cup (x_t, y_t)$; Increment t; end Return a solution: either the point evaluated with the largest f(x), or the point with the largest posterior mean.

optimization (Brochu et al., 2010b).

Bayesian optimization is also related to Bayesian active learning (or active sampling), where the goal is to efficiently learn an expensive function. The difference lies in the acquisition rule: in Bayesian optimization, acquisition rules aim at efficiently finding the maximum; whereas in Bayesian active learning, they are designed to maximize information gain about the function (see, e.g., Houlsby et al. (2011)).

3.1.2 Bayesian optimization with binary outputs

Binary Bayesian Optimization

In specific situations, direct evaluation of the objective function is impossible. This is the case when we can only measure binary outputs, such as "success" or "failure". The generalization of Bayesian optimization to these scenarios is called binary Bayesian optimization (BBO) (Tesch et al., 2013). Examples of applications include the tuning of a robot's gait parameters, where the robot may either succeed in moving over an obstacle or not (Tesch et al., 2013), or the optimization of a machine learning model's hyperparameters, by early-stopping the training of underperforming models while continuing the others (Li et al., 2018).

In BBO, for an input, x, the result of an evaluation, c, is binary (0 or 1), and assumed to be determined by a latent function, f. The output of a function follows a Bernoulli distribution with parameter π (f(x)), where π is an inverse link function used to map f(x) to [0, 1]. That is:

$$c \sim \mathcal{B}(\pi(f(\mathbf{x}))).$$

 π is typically the normal cumulative distribution function Φ . The problem is to find the maximum of *f*: that is, the input *x* such that the probability of output *c* = 1 is maximized.

Preferential Bayesian Optimization

In other cases, feedbacks are comparisons between two inputs. For example, when tuning a hearing-aid device, the user may give their preference between pairs of configurations, and the algorithm tries to learn what is the best configuration according to the user. The generalization of BO to this scenario where outputs are order relationships between pairs of inputs in so-called "duels" is thus referred to as preferential Bayesian optimization (PBO) (Gonzalez et al., 2017) or preference-based Bayesian optimization (Dewancker et al., 2017). Preference-based optimization is especially useful when the optimization process is based on human feedback. Indeed, in many situations, comparison is much easier for humans than rating (Kahneman and Tversky, 1979; Shah et al., 2014), or the value of the function to optimize is not directly measurable. This is the case in general when measuring the quality of a setting (e.g. comfort or aesthetic pleasure), as perceived by a non-expert human user. Brochu et al. (2010a,b), for example, used preference-based Bayesian optimization to set the configuration of several animation systems by querying human sequential feedback on different parameterized animations, to maximize how realistic the animation looks.

In PBO, each input x has an associated value, determined by a real valued function g to maximize, and the preference between two inputs is: f(x, x') = g(x) - g(x') (figure 3.4). The outcome of any duel (x, x') is either $x \succ x'$, or $x \prec x'$ (although PBO can be generalized to include ties, $x \approx x'$, see Dewancker et al. (2017)). The result of a duel is assumed to be stochastic, such that $(x \succ x') \sim \mathcal{B}(\pi(f(x, x')))$.

3.2 Gaussian processes as surrogate models

Bayesian optimization relies on a Bayesian surrogate model of the unknown objective function. Gaussian process (GP) models are usually chosen because of their flexibility, their practicality in incorporating prior knowledge, and because they are considered to provide a reasonable estimation of uncertainty. In this section we will present an overview of Gaussian processes, and in particular how they can be used to model binary and preference data. We also propose a minor modification to the standard Gaussian process preference learning framework and we show that it improves uncertainty estimation.

3.2.1 Gaussian processes for regression

For readers who are not familiar with stochastic processes to get an intuitive understanding, we will start with some basics and build up to Gaussian processes in their full complexity. For an excellent visual and interactive introduction to Gaussian processes, see Görtler et al. (2019).

Start by considering a Gaussian-distributed random variable Y_1 . This distribution is fully described by its mean m_1 and variance σ_1^2 . Now, consider a second variable Y_2 with marginal distribution $\mathcal{N}(m_2, \sigma_2^2)$. The Gaussian random vector (Y_1, Y_2) is described by its mean $m = (m_1, m_2)$, and its covariance matrix Σ , whose diagonal elements are σ_1^2 and σ_2^2 and elements out of diagonal are Cov (Y_1, Y_2) .

Now, assume that these two variables Y_1 and Y_2 correspond to observations at two locations in space, x_1 and x_2 . The mean and variance of Y_1 and Y_2 can thus be indexed by their

position. So we introduce functions that we call the mean function $m(\mathbf{x})$ and the covariance function $k(\mathbf{x}, \mathbf{x}')$ (or kernel) such that $m_i = m(\mathbf{x}_i)$ and $\mathbf{\Sigma} = k((\mathbf{x}_1, \mathbf{x}_2), (\mathbf{x}_1, \mathbf{x}_2))$. The variance and covariance terms are $\sigma_i^2 = k(\mathbf{x}_i, \mathbf{x}_i)$ and $\operatorname{Cov}(Y_1, Y_2) = k(\mathbf{x}_1, \mathbf{x}_2)$, respectively. Now, consider that to each location in space corresponds such a Gaussian random variable $Y(\mathbf{x})$, with mean $m(\mathbf{x})$ and variance $\sigma^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x})$. The covariance with any other set of variables at locations $X = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, denoted $Y(X) = (Y(\mathbf{x}_1), \dots, Y(\mathbf{x}_n))$, is determined by the covariance function, so that: $\operatorname{Cov}(Y(\mathbf{x}), Y(X)) = k(\mathbf{x}, X)$. This infinite collection of random variables $Y(\mathbf{x})$ is what is called a Gaussian process, that we note $\mathcal{GP}(m(\cdot), k(\cdot, \cdot))$.

Formally, a Gaussian process *Y* is a continuous space-time stochastic process such that, for any finite set of points $(x_1, ..., x_k)$, $(Y(x_1), ..., Y(x_k))$ is a multivariate Gaussian random variable.

If we sample from this Gaussian process, namely, if we jointly draw samples from these infinitely many random variables, we effectively sample a function f(x). This is why stochastic processes can be seen as distribution over functions. A GP indexed by x is fully described by a mean function mx and a covariance function k(x, x'). Moreover, Gaussian processes have the property that, if the mean and covariance functions are continuous, functions sampled from $\mathcal{GP}(m(\cdot), k(\cdot, \cdot))$ are continuous with probability one.

Since Gaussian processes are distribution over functions, they can be used as priors over functions in regression tasks (see Rasmussen and Williams (2006) for an introduction to the use of Gaussian processes in machine learning). Gaussian processes have the property of multivariate Gaussian distributions to be closed under conditioning over a finite set of observations, which makes them very easy to work with. That is, if we put a Gaussian process prior $\mathcal{GP}(m(\cdot), k(\cdot, \cdot))$ on f, then given observations $\mathbf{y} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_t)]$, the posterior over f is a GP with mean $\mu(\cdot)$, and covariance $\mathbf{\Sigma}(\cdot, \cdot)$. This property still holds with i.i.d. Gaussian noise for the observations.

Using Bayes' rule, the posterior mean μ and the posterior covariance Σ can be computed analytically to give (see appendixB):

$$\mu(\mathbf{x}) = \underbrace{m(\mathbf{x})}_{\text{prior mean}} + \underbrace{k(\mathbf{x}, X) k(X, X)^{-1} (\mathbf{y} - m(X))}_{\text{mean update}},$$
(3.2)

$$\Sigma(\mathbf{x}, \mathbf{x}') = \underbrace{k(\mathbf{x}, \mathbf{x}')}_{\text{prior covariance}} - \underbrace{k(\mathbf{x}, X) k(X, X)^{-1} k(X, \mathbf{x}')}_{\text{covariance shrinks given more data}},$$
(3.3)

with $X = (x_1, x_2, ..., x_t)$.

The computational bottleneck of Gaussian process regression is the inversion of the $t \times t$ Gram matrix k(X, X), which has complexity $O(t^3)$, where t is the number of observations. As a consequence, for applications where the amount of data is large, approximation methods have been developed (Rasmussen and Williams, 2006), or other parametric Bayesian models such as Bayesian deep neural networks can be more appropriate (Springenberg et al., 2016; Snoek et al., 2015).

Kernels and mean functions are usually parameterized using hyperparameters θ . These hyperparameters can be inferred by placing a hyperprior on them and performing Bayesian model selection (Rasmussen and Williams, 2006), by computing the posterior distribution over



Figure 3.1: Example of Gaussian process regression in a 1-dimensional input space. **A.** The red curve is the true function. We put a Gaussian process prior on this function. The blue curve is the prior mean and the shaded areas correspond to ventiles of the prior distribution. The black line is a sample from the prior distribution. **B.** The red circles correspond to data points. The blue curve is the posterior mean, and the different shaded areas correspond to ventiles of the posterior distribution. The black line is a sample from the prior distribution. **C.** Plot of the GP posterior covariance function corresponding to the regression in B. Note that since observations are noise-free, the variance collapses at observed points. Because of the smoothness prior, the variance smoothly increases as the distance from observations increases.

hyperparameters θ :

$$p(\theta|\mathscr{D}) = \frac{p(\mathscr{D}|\theta)p(\theta)}{\int p(\mathscr{D}|\theta)p(\theta)d\theta}$$
(3.4)

In practice, however, specifying an appropriate hyperprior and computing the integral in the denominator of 3.4 can be difficult. Instead, one often use type-II maximum likelihood estimation (see figure 3.2), by maximizing $p(\mathcal{D}|\theta)$ with respect to θ . Of course, care should be taken as with small amounts of data and many hyperparameters, this approach can lead to overfitting (Rasmussen and Williams, 2006).

3.2.2 Gaussian processes for classification

Gaussian processes can be used to build classification models. In binary classification, the outputs *c* can be cast into two classes, $c \in \{0,1\}$. The probability of belonging to class 1 for input *x* is assumed to be determined by a latent function *f*, such that $P(c = 1 | x) = \pi (f(x))$, where π is an inverse link function. Here we choose the normal cumulative distribution Φ as a link function. A Gaussian process prior is put on the latent function *f*.

The likelihood for this model has long been thought to be intractable and was usually computed using Laplace approximation or Expectation Propagation (Minka, 2001; Seeger, 2002) (see Rasmussen and Williams (2006) for details about these approximations and their implementation). Expectation Propagation leads to better approximations compared to Laplace approximation. However, Benavoli et al. (2020a) recently showed that the posterior can be analytically computed and corresponds to a unified skewed Gaussian process. The use of this analytical posterior led to better performance in Preferential Bayesian Optimization compared to Laplace approximation (Benavoli et al., 2020b), however, it requires a Monte-Carlo approximation of the cumulative distribution function of high dimensional multivariate Gaussians.



Figure 3.2: Gaussian process regression. The red curve is the true function sampled from a Gaussian process with a squared exponential kernel. Red circles correspond to data points. The blue curve is the posterior mean, and the different shaded areas correspond to ventiles of the predictive distribution. A. Regression obtained with random hyperparameters. In this case, the kernel lengthscale is underestimated, which explains why the posterior variance increases very rapidly as the distance from training points increases. B The same regression with hyperparameters inferred using maximum likelihood estimation. The GP extrapolates better away from data points. C. The same regression with the hyperparameters used to generate the true function. Note however, that the posterior variance with the inferred hyperparameters in B is smaller than the posterior variance with the true hyperparameters, meaning that the maximum likelihood estimation led to overestimated kernel lengthscale.

When using Laplace approximation or Expectation-Propagation (see appendix I), the posterior distribution $p(f|\mathcal{D})$ is approximated by a Gaussian process, with mean μ_f and variance σ_f^2 (see figure 3.3). The probability of output 1 for input *x* can then be expressed analytically as:

$$P(c = 1 | \mathbf{x}, \mathscr{D}) = \mathbb{E}_{f}[\Phi(f(\mathbf{x}))] = \Phi\left(\frac{\mu_{f}(\mathbf{x})}{\sqrt{1 + \sigma_{f}^{2}(\mathbf{x})}}\right)$$
(3.5)

In the remaining of this chapter, we note $\mu_c(\mathbf{x}) = P(c = 1 | \mathbf{x}, \mathcal{D})$.

3.2.3 Gaussian processes for preference learning

Pairwise preference learning as classification

Recall that preference learning aims at predicting, for any pair (x, x'), the output of a comparison between the two: either $x \succ x'$ (x is preferred over x'), or the contrary. Preference learning can be seen as a classification task, where any pair of inputs (x, x') is itself an input to a classifier with output 1 if $x \succ x'$, and 0 otherwise (Houlsby et al., 2012).

Houlsby et al. (2012), based on previous work on preference learning with Gaussian processes (Chu and Ghahramani, 2005), proposed to model preference learning as a special case of Gaussian process binary classification. The latent value function g is modelled using a Gaussian process prior with kernel $k(\cdot, \cdot)$ and prior mean $m(\cdot)$. Since the preference function f is defined as $f(\mathbf{x}, \mathbf{x}') = g(\mathbf{x}) - g(\mathbf{x}')$, this induces a GP prior on f, with prior mean $m_{\text{pref}} : \mathcal{X}^2 \to \mathbb{R}$ defined as:

$$m_{\text{pref}}\left(\boldsymbol{x}, \boldsymbol{x}'\right) = m(\boldsymbol{x}) - m\left(\boldsymbol{x}'\right), \qquad (3.6)$$



Figure 3.3: Gaussian process classification. A. The red curve is the true probability of success. Black and white dots are the data points corresponding to output 1 and 0 respectively. The blue shaded areas are the ventiles of the posterior distribution on $\Phi(f(\mathbf{x}))$, and the turquoise line is $\mu_c(\mathbf{x}) = \mathbb{E}[\Phi(f(\mathbf{x}))]$. B. The red curve corresponds to the latent function f, such that $P(c = 1|\mathbf{x}) = \Phi(f(\mathbf{x}))$. The blue curve is the posterior mean of the GP, and the shaded areas correspond to ventiles of the posterior distribution.

and a preference kernel $k_{\text{pref}} : \mathcal{X}^2 \times \mathcal{X}^2 \to \mathbb{R}$, defined as:

$$k_{\text{pref}}\left(\left(\boldsymbol{x}_{i}, \boldsymbol{x}_{j}\right), \left(\boldsymbol{x}_{k}, \boldsymbol{x}_{l}\right)\right) = k\left(\boldsymbol{x}_{i}, \boldsymbol{x}_{k}\right) + k\left(\boldsymbol{x}_{j}, \boldsymbol{x}_{l}\right) - k\left(\boldsymbol{x}_{i}, \boldsymbol{x}_{l}\right) - k\left(\boldsymbol{x}_{j}, \boldsymbol{x}_{k}\right), \quad (3.7)$$

so that $f \sim \mathcal{GP}(m_{\text{pref}}(\cdot), k_{\text{pref}}(\cdot, \cdot))$.

The result, *c*, of a comparison (x, x') is modelled as a Bernoulli random variable, such that $\mathbb{P}(c = 1 | x, x', f) = \mathbb{P}(x \succ x' | f) = \Phi(f(x, x'))$ (see figure 3.4).

Improving uncertainty estimation in Gaussian process preference learning

When using Gaussian processes for preference learning with the preference kernel, following the formalism introduced by Houlsby et al. (2012), the value function g is defined up to a constant. Indeed, all we can measure are differences between values of g at pairs of points. If the offset of g is not defined, this artificially increases the uncertainty on the preference function. This problem has not been considered in previous works.

To understand why setting the value function offset reduces uncertainty, consider that the value offset is such that: $g(x_0) = m(x_0) = 0$. The prior joint distribution over the value function *g* and the preference function *f* at points *X* is given by:

$$\begin{bmatrix} g(X) \\ g(\mathbf{x}_0) \\ f(X,\mathbf{x}_0) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} m(X) \\ m(\mathbf{x}_0) \\ m(X) - m(\mathbf{x}_0) \end{bmatrix}, \begin{bmatrix} \mathbf{K} & \mathbf{k} & \mathbf{K} - \mathbf{k} \\ \mathbf{k}^{\top} & k_0 & \mathbf{k}^{\top} - k_0 \\ \mathbf{K} - \mathbf{k}^{\top} & \mathbf{k} - k_0 & \mathbf{K}_{\text{pref}} \end{bmatrix}\right), \quad (3.8)$$

With K = k(X, X), $k = k(X, x_0)$, $k_0 = k(x_0, x_0)$ and $K_{pref} = k_{pref}((X, x_0), (X, x_0))$. This simply follows from the definition of *f* and *g*.

By applying the formula for Gaussian conditionals (B), the prior distribution for the preference function at points (X, x_0), conditioned on the value function at x_0 , is a Gaussian:

$$(f(X, \mathbf{x}_0) \mid g(\mathbf{x}_0) = 0) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \qquad (3.9)$$



Figure 3.4: Preference learning with Gaussian processes. A. Preference between \mathbf{x} and $\mathbf{x'}$ is modelled as a binary variable, with $P(\mathbf{x} \succ \mathbf{x'}) = \Phi(g(\mathbf{x}) - g(\mathbf{x'}))$, where g is the value function, \mathbf{B} . Preference can be learned from binary comparisons (black and white dots) using Gaussian process classification. White dots corresponds to $\mathbf{x} \succ \mathbf{x'}$, black dots to $\mathbf{x'} \succ \mathbf{x}$. Note that the use of a preference kernel enforces antisymmetry of the preference function, that is $f(\mathbf{x}, \mathbf{x'}) = -f(\mathbf{x'}, \mathbf{x})$. C. Posterior Gaussian process model (blue) of the value function (red). (Here, the prior GP on the value function is conditioned on $g(\mathbf{x}_0) = 0$, see section 3.2.3.)

with mean given by:

$$\mu = m_{\text{pref}} (X, \mathbf{x}_0) + (\mathbf{K} - \mathbf{k}) k_0^{-1} (g(\mathbf{x}_0) - m(\mathbf{x}_0)) = m_{\text{pref}} (X, \mathbf{x}_0), \qquad (3.10)$$

and covariance matrix given by:

$$\boldsymbol{\Sigma} = \boldsymbol{K} - \left(\boldsymbol{K} - \boldsymbol{k}\right) k_0^{-1} \left(\boldsymbol{K} - \boldsymbol{k}\right)^{\top}.$$
(3.11)

The second term in equation 3.11, $(K - k) k_0^{-1} (K - k)^\top$, corresponds to shrinkage in the predictive variance due to conditioning on $g(x_0) = 0$. This implies that setting the value function offset to some arbitrary value effectively reduces the model uncertainty. Conversely, not conditioning on this value would correspond to artificially increasing uncertainty over the preference function.

To enforce the condition $g(x_0) = 0$, one could add a virtual observation $(x_0, y_0 = 0)$ to the data. However, this would require to modify the likelihood approximation scheme. Mathematically, this is equivalent to modifying the GP prior by making use of what we call a "conditional kernel" \hat{k} :

$$\hat{k}(\boldsymbol{x},\boldsymbol{x}') = k(\boldsymbol{x},\boldsymbol{x}') - k(\boldsymbol{x},\boldsymbol{x}_0)^{\top} k(\boldsymbol{x},\boldsymbol{x}')^{-1} k(\boldsymbol{x}_0,\boldsymbol{x}')$$
(3.12)

This kernel, combined with a prior mean such that $m(x_0) = 0$, effectively enforces the condition $g(x_0) = 0$ (see the proof in appendix F). The corresponding conditional preference kernel is defined as:

$$\hat{k}_{\text{pref}}\left(\left(\boldsymbol{x}_{i},\boldsymbol{x}_{j}\right),\left(\boldsymbol{x}_{k},\boldsymbol{x}_{l}\right)\right)=\hat{k}\left(\boldsymbol{x}_{i},\boldsymbol{x}_{k}\right)+\hat{k}\left(\boldsymbol{x}_{j},\boldsymbol{x}_{l}\right)-\hat{k}\left(\boldsymbol{x}_{i},\boldsymbol{x}_{l}\right)-\hat{k}\left(\boldsymbol{x}_{j},\boldsymbol{x}_{k}\right)$$
(3.13)

Conditionally $g(\mathbf{x}_0) = 0$, the posterior distribution on g and f are linked by: $\mu_g(\mathbf{x}) = \mu_f(\mathbf{x}, \mathbf{x}_0)$ and $\sigma_g(\mathbf{x}, \mathbf{x}') = \sigma_f((\mathbf{x}, \mathbf{x}_0), (\mathbf{x}', \mathbf{x}_0))$. As a consequence, conditioning on $g(\mathbf{x}_0)$ allows to easily go back and forth from preference to value.



Figure 3.5: Comparison of posterior value GP with and without the use of the conditional preference kernel. A value function was sampled from a GP and preference data were generated randomly. The first row corresponds to a training set with 5 data points, and the second row to a training set with 100 data points. For small amount of data, conditioning on the value function offset greatly reduces posterior variance of the GP (second column, B and F), as well as $\mathbb{V}[\Phi(f(\mathbf{x}, \mathbf{x}_0))]$ (fourth column, D and H). As the amount of data increases, the difference becomes less marked. Note that the latent mean (first column, B and F) is also affected by the conditioning, just as the predictive class distribution (third column, C and G). Both are larger with conditioning. This is likely due to the coupling between between latent function variance and predictive class distribution, as $\mu_c(\mathbf{x}) = \Phi\left(\frac{\mu_g(\mathbf{x})}{\sqrt{1+\sigma_g^2(\mathbf{x})}}\right)$.

As shown on figure 3.5, the conditional preference kernel leads to dramatic predictive variance reduction for small amount of data. Note that it also induces a shift in the predictive mean. However, this effect tends to disappear as the training set grows large. To get an idea of the speed at which the effect of conditioning disappears, we measured the 2-Wasserstein distance between the latent GP with and without conditioning, averaged over 32 random functions for 50 evenly spaced test points (figure 3.6). The p-Wasserstein metric is widely used as a distance between probability distributions. The p-Wasserstein distance between the probability distributions μ and ν of two random variables X and Y is defined as:

$$W_{p}(\mu,\nu) = (\inf \mathbb{E}\left[||X - Y||^{p}\right])^{\frac{1}{p}}$$
(3.14)

Intuitively, it corresponds to the cost of optimally moving a distribution of mass μ into ν . Where the infimum is taken over all joints distributions of *X* and *Y* with marginals μ and ν . For multivariate Gaussian distributions, the 2-Wasserstein distance W_2 is given by:

$$W_{2}\left(\mathcal{N}(\boldsymbol{m}_{1},\boldsymbol{K}_{1}),\mathcal{N}(\boldsymbol{m}_{2},\boldsymbol{K}_{2})\right)^{2} = \|\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\|^{2} + \operatorname{tr}\left(\boldsymbol{K}_{1}+\boldsymbol{K}_{2}-2\left(\boldsymbol{K}_{1}^{\frac{1}{2}}\boldsymbol{K}_{2}\boldsymbol{K}_{1}^{\frac{1}{2}}\right)^{\frac{1}{2}}\right)$$
(3.15)

 m_i and K_i are the posterior mean and covariance of each distribution. On figure 3.6, we see that, as the number of training points increase, the 2-Wasserstein distance between the Gaussian process conditioned on $g(x_0 = 0$ and the unconditioned one rapidly decrease so that the



Figure 3.6: Effect of value function conditioning on preference learning, as a function of training set size. We measured the 2-Wasserstein distance between the GP conditioned on $g(x_0) = 0$ and the GP that was not conditioned. The distance diminishes as the training set size increases and becomes negligible.

difference in the two distribution becomes negligible beyond about 100 training points. This can be interpreted as the prior becoming less and less influential on the posterior. However, in Bayesian optimization, one often deals with very small dataset, so that this conditioning could make a difference.

3.2.4 Kernel and mean

The mean and covariance functions can be used to incorporate prior information about the task at hand. When no information is available, a common choice is to set the mean to a constant value. A Gaussian process model acts as an interpolator, and the kernel determines the dependency between two data points. It encodes assumptions about the function such as smoothness, periodicity, symmetries, and invariances (see Duvenaud (2014) for a deep dive into the design and choice of kernels).

Kernels are often chosen to be stationary, i.e. k(x, x') = k(x - x'). For GP priors with constant mean, this encodes the prior belief that the difference between function values is invariant to translations. Figure 3.7 shows examples of kernels and samples from the corresponding priors.

Since the kernel directly determines the surrogate model in Bayesian optimization, the choice of kernel's hyperparameters can have a dramatic effect on the performance of Bayesian optimization, as shown in figure 3.8.

3.3 Acquisition rules with binary outputs

So far, we have considered the model used to describe the available knowledge about the objective function. In Bayesian optimization, this model is used to guide the process of selecting new observations so as to efficiently find the maximum of the objective, and eventually output a solution. The challenge of designing acquisition rules is to balance exploration and exploitation in the search space so as to efficiently optimize the objective.



Figure 3.7: Upper row: Examples of kernels and the corresponding covariance matrices. Bottom row: Gaussian processes with the different kernels and zero mean (blue), and samples from these distributions (black). The shaded regions correspond to ventiles of the distribution at the corresponding location.

3.3.1 Acquisition in standard Bayesian optimization

At iteration *t*, a utility $U(x, y | \mathcal{D}_t)$ of the next observation (x, y) can be defined based on the surrogate model. This utility is somehow related to how much observation (x, y) improves our solution to the objective maximization problem 3.1. The maximum of the expected utility, $\alpha(x) = \mathbb{E}_y(U(x, y) | \mathcal{D}_t)$, is called an acquisition function, and is a natural choice for the point *x* where to evaluate the objective. However, we will see in subsequent sections that acquisition rules exist that actually deviate from this formalism. In a fully Bayesian treatment, expectation should also be taken over the kernel and mean hyperparameters θ .

Acquisition rules are very diverse and reflect different heuristics used to tradeoff exploration and exploitation. Exploration is the selection of queries for which the output is highly uncertain. A purely exploratory step is likely not to improve the solution to the optimization problem. Exploitation, on the other hand, is the selection of the point minimizing the expected instantaneous regret. Exploitation tend to focus on regions where uncertainty is small, and where the maximum is likely to lie. As a consequence, exploration and exploitation are usually antagonistic.

In standard Bayesian optimization, where evaluation of the function is noise-free, the most commonly used rule is to select the point that maximizes expected improvement (Mockus, 1989; Jones et al., 1998). For an input x, improvement corresponds to how much a new function evaluation in x improves our solution to problem 3.1. For noise-free observations, the reported solution if we would have to stop evaluating the function is $f^{**} = \max_{x \in x_{1:t}} f(x)$, the maximum among previous observations (we use the double star notation to emphasize that the maximum



Figure 3.8: Effect of kernel hyperparameters on the performance of Bayesian optimization. Left panel: A function was sampled from a six-dimensional Gaussian process (projection on the first axis in blue) and used as an objective function. Changing the hyperparameters changes the smoothness of the samples (red). After learning the hyperparameters with maximum-likelihood inference based on 50 observations, we recover the correct behavior (green). Right panel: the blue function was optimized using Expected Improvement with the correct hyperparameters used to learn the surrogate model (blue), the wrong hyperparameters (red), or while learning the hyperparameters using maximum-likelihood at each iteration (green). The first five queries were selected at random. The lines correspond to the average best value observed across 20 repetitions of the experiment, whereas the shaded area corresponds to standard error on the mean. With the wrong hyperparameters, the optimization consistently fails, whereas learning the hyperparameters slows the optimization down, but eventually leads to finding good solutions.

is taken among previous observations).

For a new observation f(x), improvement is thus defined as:

$$I(\mathbf{x}) = [f(\mathbf{x}) - f^{\star\star}]_{+}$$
(3.16)

As a consequence, expected improvement is defined as:

$$\alpha_{EI}(\mathbf{x}) = \mathbb{E}\left[I(\mathbf{x})|\mathscr{D}_t\right] = \mathbb{E}\left[\left[f(\mathbf{x}) - f^{\star\star}\right]_+ |\mathscr{D}_t\right]$$
(3.17)

Expected improvement can be evaluated analytically:

$$\alpha_{EI}(\mathbf{x}) = \int_{I=0}^{I=\infty} \frac{I}{\sqrt{2\pi}\sigma(\mathbf{x})} \exp\left(-\frac{(\mu(\mathbf{x}) - f^{\star\star} - I)^2}{2\sigma^2(\mathbf{x})}\right) dI$$

= $(\mu(\mathbf{x}) - f^{\star\star}) \Phi\left(\frac{\mu(\mathbf{x}) - f^{\star\star}}{\sigma(\mathbf{x})}\right) + \sigma(\mathbf{x})\phi\left(\frac{\mu(\mathbf{x}) - f^{\star\star}}{\sigma(\mathbf{x})}\right)$ (3.18)

Expected improvement performs well compared to random sampling when measuring the cumulative regret on synthetic problems (figure 3.9.C). One of the key properties of the expected improvement acquisition function is that it increases with respect to both the posterior mean and variance, meaning that it favors points whose value is either uncertain (exploration) or believed to be high (exploitation) (Frazier, 2018).

3.3.2 Exploration and exploitation in BO with binary outputs

When outputs are binary, specific difficulties arise that do not occur in standard Bayesian optimization because the underlying objective is not directly measured. To design acquisition rules 3.3. Acquisition rules with binary outputs



Figure 3.9: Bayesian optimization with continuous outputs. A. Red: true function. Red circles: data points. Blue: Posterior GP. **B.** Expected improvement: The next sample is the one maximizing expected improvement. Thompson sampling: The next query is a sample from $p(\mathbf{x}^*|\mathcal{D}_t)$. In practice, a sample is drawn from p(f|D) and is then maximized. **C.** Cumulative regret averaged over 300 repetitions of the same optimization problem. The hyperparameters are assumed to be known. The shaded area width corresponds to the standard error on the mean. The cumulative regret with expected improvement saturates, meaning that the optimization succeeds. On the contrary, the random acquisition leads to a much slower decrease of regret.

in the case of binary outputs, a first step is to understand what exploration and exploitation mean in this context.

Exploitation

An exploitative strategy is defined as selecting the location minimizing the expected instantaneous regret. Given an identification procedure that returns x_t^* at time t, a natural definition of the instantaneous regret at time t in standard BO is $R_t = f(x^*) - f(x_t^*)$, where x^* is the true maximum of f. In standard BO, the query minimizing the expected regret is thus simply arg max f(x). The corresponding acquisition rule is called Simple Regret (see figure 3.11). $x \in x_{1:t}$

In the case of binary BO, since we are interested in maximizing the probability of success, the non-linearity has to be taken into account. Pure exploitation would correspond to selecting $\underset{x \in \mathcal{X}}{\arg \max \mu_c(x)}$. Finally, in PBO, pure exploitation would correspond to selecting as first and second duel members $\arg \max \mu_g(x)$, so it is particularly clear that pure exploration does not make sense in this case.

Exploration and epistemic uncertainty

Pure exploration can be seen as sampling at locations where uncertainty is maximal. In general, uncertainty about the outcome of an observation has two components: an aleatoric one, which stems from the fact that the outputs of a function evaluation are intrinsically random; and an epistemic component, coming from the uncertainty about the underlying function. This distinction is not important in standard Bayesian optimization, without noise or with homoscedastic noise. Indeed, in BO with noise-free observations, uncertainty has no aleatoric component. With i.i.d. Gaussian noise, aleatoric uncertainty is constant. In GP regression, pure exploration

thus corresponds to selecting the location x with maximum GP posterior variance $\sigma_f^2(x)$. However, this distinction between the two components of uncertainties becomes crucial in situations where aleatoric uncertainty varies with function values. For example, in GP classification, even when the latent function f is fully known, there is uncertainty remaining about the outcome c, and this uncertainty depends on the value of f because the probability that the output c = 1 is given by $\Phi(f(x))$.

To see why the distinction between the two components of uncertainty is important, consider the case of BBO, where selecting an input x with very little posterior variance of f(x) but posterior mean equal to zero will result in $p(c = 1 | x, \mathcal{D}_t) = \frac{1}{2}$. As a consequence, even though f(x) is known with high precision, the output of the query is maximally uncertain. On the contrary, selecting a point with high posterior variance of f(x), but large posterior mean would result in $p(c = 1 | x, \mathcal{D}_t) = 1$. In both cases, the observation is expected to bring very little information about the underlying function.

In this context, if the variance is used as a measure of uncertainty, the predictive variance $\mathbb{V}(c|\mathcal{D}_t, x)$ corresponds to the total uncertainty at point x. In Bayesian learning, both aleatoric and epistemic uncertainty are modelled using the same probability measure, but they can be distinguished by applying the law of total variance:

$$\mathbb{V}(c|\mathscr{D}_t, \mathbf{x}) = \mathbb{V}_{p(f|\mathscr{D}_t)}[\Phi(f(\mathbf{x}))] + \mathbb{E}_{p(f|\mathscr{D}_t)}[\mathbb{V}(c|f, \mathbf{x})]$$

Depeweg et al. (2017) interpreted this equation (in the case of Bayesian neural networks) as a decomposition into aleatoric and epistemic uncertainty. $\mathbb{E}_{p(f|\mathscr{D}_t)}[\mathbb{V}(c|f, \mathbf{x})]$ corresponds to the average value of the variance of *c* conditioned on *f*, and thus represents aleatoric uncertainty. Indeed, consider the case where *f* is fully known, then the epistemic uncertainty is removed, so $\mathbb{V}(c|f, \mathbf{x})$ is a natural measure of aleatoric uncertainty. Since *f* is not known, aleatoric uncertainty is measured in terms of expectation with respect to $p(f|\mathscr{D}_t)$ (Hossein Shaker and Hüllermeier, 2020).

The aleatoric uncertainty term can be expressed as:

$$\mathbb{E}_{p(f|\mathscr{D}_t)}[\mathbb{V}(c|f, \mathbf{x})] = \mathbb{E}_{p(f|\mathscr{D}_t)}[\Phi(f(\mathbf{x}))(1 - \Phi(f(\mathbf{x})))]$$

The term $\mathbb{V}_{p(f|\mathscr{D}_t)}[\mathbb{E}(c|f, \mathbf{x})] = \mathbb{V}_{p(f|\mathscr{D}_t)}[\Phi(f(\mathbf{x}))]$ corresponds to epistemic uncertainty, that is, the uncertainty about *c* that comes from not knowing *f*. We will see in section 3.3.2 that it can be analytically computed.

A natural definition of exploration is to select points where epistemic uncertainty is maximal. As noted by Gonzalez et al. (2017) who first proposed this definition of exploration in the context of PBO, this corresponds to selecting the point where the Gaussian process is the most uncertain about the probability of the outcome $\Phi(f(\mathbf{x}))$.

Another approach is to express uncertainty using entropy. The total uncertainty about the output of a query is $H(c|\mathscr{D}_t, \mathbf{x})$. Aleatoric uncertainty, on the other hand, is the uncertainty about the output when the underlying function is known, averaged over possible latent function values, that is $\mathbb{E}_{p(f|\mathscr{D}_t)}[H(c|f, \mathbf{x})]$. The difference between these two quantities is epistemic uncertainty. However, $H(c|\mathscr{D}_t, \mathbf{x}) - \mathbb{E}_{p(f|\mathscr{D}_t)}[H(c|f, \mathbf{x})]$ is also the mutual information between the output and the latent function: $I(c, f|\mathbf{x}, \mathscr{D}_t)$. Depeweg et al. (2017) thus proposed the fol-

lowing decomposition (see figure 3.10):

$$\underbrace{H(c|\mathcal{D}_t, \mathbf{x})}_{\text{Total uncertainty}} = \underbrace{I(c, f|\mathcal{D}_t, \mathbf{x})}_{\text{Epistemic uncertainty}} + \underbrace{\mathbb{E}_{p(f|\mathcal{D}_t)}[H(c|f, \mathbf{x})]}_{\text{Aleatoric uncertainty}}$$

A priori, it may seem counter-intuitive to define epistemic uncertainty using mutual information, since the more information we have access to, the least uncertain we are. However, if we note that mutual information between c and f corresponds to how much knowing f reduces the uncertainty about c, the equivalence between mutual information and epistemic uncertainty about c becomes intuitive.

When making an observation at a point *x*, the mutual information $I(c, f | \mathcal{D}_t, x)$ is the expected decrease in entropy of *f* that is due to observing *c*:

$$I(c, f|\mathscr{D}_t, \mathbf{x}) = H(f|\mathscr{D}_t, \mathbf{x}) - \mathbb{E}_{p(c|\mathscr{D}_t, \mathbf{x})}[H(f|c, \mathbf{x})]$$
(3.19)

As a consequence, selecting the point maximizing epistemic uncertainty (known as epistemic uncertainty sampling (Nguyen et al., 2019)) corresponds to selecting the point maximizing the expected decrease in latent function entropy. This information-theoretic acquisition criterion was introduced in Bayesian active learning by Houlsby et al. (2011), under the name Bayesian Active Learning by Disagreement (BALD), because it corresponds to selecting the point *x* for which the posterior on *f* disagree about the outcome *c* the most. In GP classification and preference learning, $I(c, f | x, \mathcal{D}_t)$ can be approximated and maximized at a low cost (Houlsby et al., 2011). Figure 3.11 shows a comparison of different exploration strategies in a classification task.

As we will see in 3.3.5, the lack of proper definition of exploration led to a significant confusion in the design of acquisition rules, with some authors considering for example that maximizing the latent variance corresponded to exploration. A mistake, for example, is to define exploration in the latent space, by considering that exploration corresponds to sampling at locations of high posterior variance of f, independently of the value of the posterior mean. As can be seen of figure 3.10, even with large posterior variance of the latent GP, total uncertainty can be small provided the posterior mean is large (or small) enough.

Epistemic uncertainty of a Gaussian process classification model

In the paper by Gonzalez et al. (2017), epistemic uncertainty $\mathbb{V}[\Phi(f(\mathbf{x}))]$ was computed using a Monte-Carlo approximation. However, we found that an analytic form can be derived, with a tractable gradient. Indeed, by definition:

$$\mathbb{V}\left[\Phi(f(\mathbf{x}))\right] = \int \left(\Phi(f(\mathbf{x})) - \mathbb{E}\left[\Phi(f(\mathbf{x}))\right]\right)^2 p(f \mid \mathscr{D}_t, \mathbf{x}) df$$
(3.20)

As a consequence:

$$\begin{aligned} \mathbb{V}\left[\Phi(f(\mathbf{x}))\right] &= \int \Phi(f(\mathbf{x}))^2 p(f(\mathbf{x}) \mid \mathscr{D}_t, \mathbf{x}) df - \mathbb{E}\left[\Phi(f(\mathbf{x}))\right]^2 \\ &= \int \Phi\left(\sigma_f(\mathbf{x}) f(\mathbf{x}) + \mu_f(\mathbf{x})\right)^2 \phi\left(f(\mathbf{x})\right) df - \mathbb{E}\left[\Phi\left(f(\mathbf{x})\right)\right]^2, \end{aligned} \tag{3.21}$$

where ϕ is the probability density function of $\mathcal{N}(0,1)$, μ_f and σ_f^2 are the mean and variance of the Gaussian approximation to the posterior $p(f|\mathcal{D}_t)$. Making use of the following formula





Figure 3.10: Uncertainty decomposition in Gaussian process classification. The horizontal and vertical axes correspond to the posterior mean $\mu_f(\mathbf{x})$ and variance $\sigma_f^2(\mathbf{x})$ of the latent Gaussian process. The first row corresponds to uncertainty expressed in terms of variance, whereas the second row corresponds to the information-theoretic decomposition. *A*,*D*. Total uncertainty. *B*,*E*. Epistemic uncertainty, measured as the mutual information between the output c and $f(\mathbf{x})$ (B) or as the posterior variance of $\Phi(f(\mathbf{x}))$ (E). *C*,*F*. Aleatoric uncertainty, measured as $\mathbb{E}_{p(f|\mathscr{D}_t)}[H(c|f,\mathbf{x})]$ (C) or the expected variance of the output c given f. Note that both epistemic and aleatoric uncertainty increase as μ_f gets closer to 0, but aleatoric uncertainty decreases with σ_f^2 whereas epistemic uncertainty increases. Epistemic uncertainty is related to how informative a function evaluation is expected to be. In GP classification, contrary to standard GP regression, epistemic uncertainty depends on the latent mean

(Patel and Read, 1996):

$$\int_{-\infty}^{\infty} \Phi(a+bt)^2 \phi(t) dt = \Phi\left(\frac{a}{\sqrt{1+b^2}}\right) - 2T\left(\frac{a}{\sqrt{1+b^2}}, \frac{1}{\sqrt{1+2b^2}}\right),$$
 (3.22)

with the Owen's T function defined as:

$$T(h,a) = \frac{1}{2\pi} \int_0^a \frac{e^{-\frac{1}{2}h^2(1+t^2)}}{1+t^2} dt,$$
(3.23)

equation 3.20 can be written as:

$$\mathbb{V}\left[\Phi\left(f(\boldsymbol{x})\right)\right] = \mathbb{E}\left[\Phi\left(f(\boldsymbol{x})\right)\right] - 2T\left(\frac{\mu_f(\boldsymbol{x})}{\sqrt{1 + \sigma_f^2(\boldsymbol{x})}}, \frac{1}{\sqrt{1 + 2\sigma_f^2(\boldsymbol{x})}}\right) - \mathbb{E}\left[\Phi\left(f(\boldsymbol{x})\right)\right]^2.$$
(3.24)

The Owen's T function can be cheaply evaluated numerically (Young and Minder, 1974; Thomas, 1986). The computation of the gradient of $\mathbb{V}[\Phi(f(x))]$ with respect to x is thus straightforward (see appendix D).

The component of variance corresponding to aleatoric uncertainty is:

$$\mathbb{E}[\mathbb{V}(c|f, \mathbf{x})] = \mathbb{E}[\Phi(f(\mathbf{x}))(1 - \Phi(f(\mathbf{x}))] \\ = \mathbb{E}\left[\Phi(f(\mathbf{x}))\right] - \mathbb{E}\left[\Phi(f(\mathbf{x}))^2\right]$$
(3.25)



Figure 3.11: Active learning. A. GP classification model. The latent function is linear. The true class probability distribution is plotted in red, the predictive class distribution is plotted in turquoise and the posterior distribution on $\Phi(f(\mathbf{x}))$ is in blue (shaded areas correspond to ventiles of the distribution). **B**. Exploration by maximizing variance (in red) and Bayesian Active Learning by Disagreement (BALD, Houlsby et al. (2011)) (in blue). **C**. Comparison of three active learning strategies on 20 repetitions of synthetic active learning experiments on the function plotted on panel A. The regret is the root mean squared error between P(y = 1) and μ_c . In this example, the two strategies that actively explore by selecting points maximizing epistemic uncertainty are as efficient, and significantly fasten learning compared to random selection.

Moreover, we note that epistemic uncertainty can be written as:

$$\mathbb{V}\left[\Phi\left(f(\boldsymbol{x})\right)\right] = \mathbb{E}\left[\Phi\left(f(\boldsymbol{x})\right)^{2}\right] - \mathbb{E}\left[\Phi\left(f(\boldsymbol{x})\right)\right]^{2}$$
(3.26)

By combining equations 3.25 and 3.26 with result 3.24, we see that the aleatoric uncertainty component is:

$$\mathbb{E}[\mathbb{V}(c|f, \mathbf{x})] = 2T\left(\frac{\mu_f(\mathbf{x})}{\sqrt{1 + \sigma_f^2(\mathbf{x})}}, \frac{1}{\sqrt{1 + 2\sigma_f^2(\mathbf{x})}}\right)$$
(3.27)

3.3.3 Selecting the best point after spending the evaluation budget

The choice of an acquisition rule depends on the method used to report a solution to the maximization problem. Indeed, after spending the evaluation budget *T*, the algorithm has to output a solution in what is called an identification step (Jalali et al., 2017). In standard Bayesian optimization, without noise, one possibility is to select the point among the previous observations with the highest value. Of course, other possibilities exist, such as selecting the point with the highest predictive mean. However, to limit risk in practice, one may want to select a solution among points that have been evaluated.

In the noisy setting, the choice of the identification procedure is not trivial. A typical strategy is to choose the maximum of the posterior mean, independent of the variance (see figure 3.12). Another possibility would be to select the most likely maximum, that is the maximum of $p(\mathbf{x}^*|\mathcal{D}_t)$ (Hennig and Schuler, 2012). However, this requires building a Monte Carlo approximation of $p(\mathbf{x}^*|\mathcal{D}_t)$.

In Bayesian optimization with binary outcomes, the problem of choosing an identification



Chapter 3. Practical Bayesian optimization with binary and preference data

Figure 3.12: Identification step. Red: true function. Red circles: data points, evaluation of the function is noisy. Blue: Posterior GP. Note that the kernel of the GP takes the i.i.d. Gaussian noise into account. Black: approximate $p(\mathbf{x}^* | \mathscr{D}_t)$, obtained by drawing 50000 samples from $p(f | \mathscr{D}_t)$ and maximizing them. The behavior of the distribution at the boundaries of the domain is due to the smoothness prior. Indeed, for a point on the left (resp. right) of the boundary, the derivative of f is positive (resp. negative) with non-zero probability. So with non-zero probability, the maximum lies exactly on the boundaries of the domain (Hennig and Schuler, 2012). The blue dotted line corresponds to the maximum of the posterior mean. The black dotted line corresponds to the maximum of $p(\mathbf{x}^* | \mathscr{D}_t)$. Here, the maximum of the posterior mean coincides with the maximum among previously evaluated points.

procedure is similar. For binary Bayesian optimization , Tesch et al. (2013) choose the maximum of the posterior probability of success, i.e. arg max $\mu_c(\mathbf{x})$.

 $x \in \mathcal{X}$

For preferential Bayesian optimization, Gonzalez et al. (2017) suggested to use a different identification step. They applied a modification of the concept of Copeland score, used in the ranking literature (Zoghi et al., 2015). The normalized Copeland score is defined as:

$$S(\mathbf{x}) = \operatorname{Vol}(\mathcal{X})^{-1} \int_{\mathcal{X}} \mathbb{I}_{\{\Phi[f(\mathbf{x}, \mathbf{x}')] \ge 0.5\}} d\mathbf{x}', \qquad (3.28)$$

where $\operatorname{Vol}(\mathcal{X}) = \int_{\mathcal{X}} d\mathbf{x}'$ is a normalizing constant. $S(\mathbf{x})$ is the proportion of duels that \mathbf{x} will most likely win. Since the maximum \mathbf{x}^* of g is the only \mathbf{x} for which $\forall \mathbf{x}' \in \mathcal{X}, \Phi[f(\mathbf{x}, \mathbf{x}')] \ge 0.5$, we have $S(\mathbf{x}^*) = 1$, so that \mathbf{x}^* is the maximum of S.

Gonzalez et al. (2017) defined the soft-Copeland score as:

$$C(\mathbf{x}) = \operatorname{Vol}(\mathcal{X})^{-1} \int_{\mathcal{X}} \Phi[f(\mathbf{x}, \mathbf{x}')] d\mathbf{x}'$$
(3.29)

This corresponds to the average probability of winning a duel. The Condorcet winner at time *t* is defined as the maximum of the soft-Copeland score:

$$\mathbf{x}_{c,t} = \arg\max_{\mathbf{x}\in\mathcal{X}} C(\mathbf{x}) \tag{3.30}$$

The Condorcet winner is the point that has, on average, the maximum probability of winning duels, and therefore, it is, again, the optimum of *g*.

The expectation of the soft-Copeland score is ¹:

¹In the paper, the authors indistinctly used C(x) and $\mathbb{E}[C(x)]$.

$$\mathbb{E}[C(\mathbf{x})] = \operatorname{Vol}(\mathcal{X})^{-1} \int_{\mathcal{X}} \mu_c(\mathbf{x}, \mathbf{x}') d\mathbf{x}'$$
(3.31)

Gonzalez et al. (2017) choose to return the maximum of $\mathbb{E}[C(\mathbf{x})]$ as the solution. However, this identification step is difficult to use in practice as it requires to estimate an intractable integral, with cost that grows exponentially with the dimension of \mathcal{X} . In practice, most authors simply select the maximum of the latent mean μ_g as the solution. This choice is justified by the fact $\mathbf{x}_t^* = \underset{x \in \mathcal{X}}{\operatorname{arg max}} \mu_g(\mathbf{x})$ is the most likely winner of any duel. To see this, note that for any duel $(\mathbf{x}_1, \mathbf{x}_2)$:

$$P(\mathbf{x}_1 \succ \mathbf{x}_2 | \mathscr{D}_t) = \Phi\left(\frac{\mu_g(\mathbf{x}_1) - \mu_g(\mathbf{x}_2)}{\sqrt{1 + \sigma_f^2(\mathbf{x}_1, \mathbf{x}_2)}}\right)$$
(3.32)

So, for all x, $P(x_t^* \succ x) \ge \frac{1}{2}$.

To understand how the identification step is linked to the choice of an acquisition rule, consider for example expected improvement. The notion of improvement depends itself on the way performance is measured, that is, on the identification procedure. As a consequence, implicit in many acquisition rule is the identification step.

3.3.4 Acquisition rules for binary BO

Armed with a clear definition of the exploration/exploitation tradeoff, we can consider the generalization of acquisition functions used in standard BO to the binary case.

Binary Expected Improvement The generalization of EI to the noisy or binary outcomes scenario is not trivial, as the notion of improvement is not easily defined. This problem is discussed in the noisy setting in section 5 of Frazier (2018) and Letham et al. (2019). Various improvement-based heuristics have been proposed that use different definitions of improvement.

Given an identification step in which the reported solution is the maximum of the posterior mean of the latent GP among previously sampled points, the natural generalization of Expected Improvement in the situation where outputs are binary is: $\mathbb{E}(\mu_{t+1}^{\star\star} - \mu_t^{\star\star} | x_{t+1})$, where the expectation is taken over the possible outcomes, and where $\mu_t^{\star\star} = \max_{x \in X_t} \mu_t(x)$. However, given that we do not have direct access to f, there is no reason to restrict ourselves to previous observations. Tesch et al. (2013) proposed to use as reported solution the maximum of $\mu_c(x)$, and introduced the binary expected improvement:

$$\alpha_{\text{BEI}}(\mathbf{x}) = \mathbb{E}\left[\left[\Phi(f(\mathbf{x})) - \max_{\mathbf{x} \in \mathcal{X}}(\mu_c(\mathbf{x}))\right]_+\right]$$
(3.33)

However, this acquisition function is intractable and has to be estimated via numerical integration. However, this integration does not depend on the dimension of the input space and is thus relatively cheap.



Figure 3.13: Comparison of acquisition functions for binary BO. The min-max normalized values of different acquisition function is plotted as a function of the latent mean $\mu_f(\mathbf{x})$ and variance $\sigma_f^2(\mathbf{x})$. Expected improvement of in the latent space (EI_f, computed assuming $\mu(\mathbf{x}^{\star\star}) = 0$) favors exploitation by favoring points with high posterior mean. Expected improvement in the probability space EI_Φ (computed using Monte-Carlo integration, assuming $\mu(\mathbf{x}^{\star\star}) = 0$) is more explorative than EI_f. UCB_f explicitly balances exploration and exploitation in the latent space, favoring points with high posterior mean, independently of the fact that these points are expected to bring very little information about the underlying function. This problem is solved in UCB_Φ, which trades off exploration and exploitation in the probability space. With UCB_Φ, the expected utility does not monotonically increases with the posterior latent mean $\mu_f(\mathbf{x})$.

Upper Credible Bound Tesch et al. (2013) compared this acquisition rule to EI on the latent function and to Upper Credible Bound (UCB) in the latent space (UCB_{*f*}), defined as:

$$\alpha_{\rm UCB_f}(\mathbf{x}) = \mu_f(\mathbf{x}) + \beta_t \sigma_f(\mathbf{x})$$

Where β_t are constants that determine the tradeoff between exploration and exploitation. Higher values of β_t favors more exploratory behavior. BEI outperformed acquisition rules that work in the latent space, probably because they under-explore for the reasons explained in section 3.3.2. For example, UCB_f tend to select points with high latent mean (figure 3.13), despite the fact that for large posterior variance the output at these locations is very likely to be c = 1, and observations will bring little information (see figure 3.14 for a pathological example).

As a consequence, it is tempting to generalize the very popular UCB acquisition function by working in the space of probabilities. Based on our discussion on exploration in Bayesian optimization, we thus propose the following generalization of UCB to binary BO:

$$\alpha_{\text{UCB}_{\Phi}}(\mathbf{x}) = \mathbb{E}[\Phi(f(\mathbf{x})] + \beta \sqrt{\mathbb{V}[\Phi(f(\mathbf{x}))]}$$
(3.34)

The constant β determines the tradeoff between exploration (high epistemic uncertainty $\mathbb{V}[\Phi(f(\mathbf{x}))]$) and exploitation, that is, high probability of c = 1, given by $\mathbb{E}[\Phi(f(\mathbf{x}))]$ (we choose $\beta = \Phi^{-1}(0.99)$). Contrary to Binary Expected Improvement, this acquisition has an analytical expression (see equation 3.24) as well as its gradient (see appendix D). The difference in behavior is illustrated on figure 3.13.

Binary Knowledge Gradient If we were to report our solution after the next query, then a natural choice would be to select the observation that is expected to maximally improve its



Figure 3.14: Comparison between UCB in the latent space (UCB_f) and UCB in the probability space (UCB_Φ). A. Probability of output c = 1 as a function of input \mathbf{x} (in red), posterior distribution $p(\Phi(f(\mathbf{x}))|\mathcal{D})$ (in blue, shaded areas correspond to ventiles of the distribution), and predictive class distribution $\mu_c(\mathbf{x} \text{ (in turquoise)})$. The choice of the next according to UCB_Φ and UCB_f is indicated by crosses. B. True latent function f (in red) and posterior Gaussian process $p(f|\mathcal{D})$ (in blue). C. Upper Credible Bound acquisition function in the latent space. The selected point is indicated with a dashed line. D. Upper Credible Bound acquisition function function in the latent space. The selected point is indicated with a capuisition function tends to select points with high posterior latent mean, despite the fact that these points correspond to very little uncertainty about the output. This under-exploration problem does not appear with UCB_Φ.
value. If we choose to report as a solution: $x_{t+1}^* = \underset{x \in \mathcal{X}}{\operatorname{arg\,max}} \mathbb{E}(f(x)|\mathscr{D}_{t+1})$, then with $\mu_t^* = \underset{x \in \mathcal{X}}{\operatorname{max}} \mathbb{E}(f(x)|\mathscr{D}_t)$, $\mu_{t+1}^* - \mu_t^*$ corresponds to the change in the expected value of the solution after observing (x_{t+1}, c_{t+1}). This is the rationale behind the knowledge gradient acquisition rule (Frazier et al., 2009), defined as:

$$\alpha_{\mathrm{KG}}(\mathbf{x}) = \mathbb{E}(\mu_{t+1}^{\star} - \mu_t^{\star} | \mathcal{D}_t, \mathbf{x}_{t+1} = \mathbf{x})$$
(3.35)

Here, we propose to generalize the knowledge gradient to the binary feedback scenario with Gaussian process classification models. For efficient optimization, we also derived an approximation of the gradient of the knowledge gradient in the case of the Laplace approximation (see appendix E). Wang et al. (2016) introduced the knowledge gradient acquisition rule for the same class of problems (KG), but with a linear Bayesian classifier as a surrogate model. They showed that when using a linear Bayesian classifier, the binary knowledge gradient outperforms expected improvement, as well as UCB_f and Thompson sampling (TS). Thompson sampling (Thompson, 1933) is a general strategy that consists in choosing the action which maximizes the expected reward compared to a belief drawn at random. In the case of Bayesian optimization, it consists in selecting an observation at a sample from the posterior probability distribution of the maximum $p(\mathbf{x}^* | \mathcal{D}_t)$. This sampling comes with several difficulties, which we describe and address in detail in section 3.6.

The various acquisition functions for binary Bayesian optimization are summarized in table 3.1. To illustrate the behavior of the different acquisition rules, one can visualize the Pareto front of the multi-objective maximization problem corresponding to maximizing exploitation and exploration (figure 3.15):

$$\underset{\mathbf{x}\in\mathcal{X}}{\arg\max}\left(\mathbb{E}[\Phi(f(\mathbf{x}))|\mathcal{D}],\sqrt{\mathbb{V}[\Phi(f(\mathbf{x}))|\mathcal{D}]}\right).$$
(3.36)

A one-step ahead optimal rule should select a point of the Pareto front, as otherwise a more exploratory and exploitative choice could be made. By definition, UCB_{Φ} selects optimal points (with the position on the front determined by β). In general, expected improvement tend to be very exploitative. Thompson sampling is not guaranteed to select points on the Pareto front. Note, however, that it does not mean that Thompson sampling is a suboptimal rule. Indeed, non-myopic rules may be for example two-step ahead optimal without being one-step ahead optimal.

3.3.5 Acquisition rules for PBO

The main difference between BO and Preferential Bayesian Optimisation is that, in the latter, the acquisition function returns a new duel instead of a single input. We will now describe the various heuristics that have been proposed.

The different heuristics used in PBO can be classified according to whether the two duel members are selected jointly or sequentially. The sequential rules can be further classified by considering how the first duel member is selected. In most applications of PBO, the first duel member is either the maximum of the posterior mean of the latent function among the training points or the maximum of the predictive mean (denoted as the "champion" in the dueling bandits literature (Bengs et al., 2021)). The second class of rules selects the first duel member



Figure 3.15: Example of acquisition rules behavior in BBO. A. The posterior distribution $P(c = 1|\mathscr{D})$ is drawn in blue, shaded areas correspond to ventiles of the distribution. The turquoise line is the mean of the predictive class distribution, $\mathbb{E}[\Phi(f(\mathbf{x}))|\mathscr{D}]$. The gray line is the true distribution, and superimposed in red is the Pareto front of the maximization problem $\max_{x \in \mathcal{X}} (\mathbb{E}[\Phi(f(\mathbf{x}))|\mathscr{D}], \mathbb{V}[\Phi(f(\mathbf{x}))|\mathscr{D}])$. The points selected by the different acquisition rules are indicated with crosses. **B.** Plot of the Pareto front corresponding to the Gaussian process in *A*, plotted in the 'mean-variance space'. The Pareto front is drawn in red, and the Pareto-dominated points are drawn in gray. The points selected by the different acquisition sample, all acquisition rules except UCB_{\Phi} very strongly exploit. Note also that both Thompson sampling and Expected improvement select points that are Pareto-dominated, meaning that they do not correspond to an optimal trade-off between exploration and exploitation. This representation is inspired by De Ath et al. (2019).

using Thompson Sampling. The problem is then to select a "challenger". The second member is usually chosen according to heuristics from standard Bayesian optimization (the behavior of various heuristics used to select challengers is illustrated in figure 3.18). Benavoli et al. (2020b), for example, proposed three acquisition rules which all take as the first duel member x_1 the maximum of the posterior mean of the latent value function, and select x_2 either by sampling $p(x^*|\mathcal{D}_t)$ (dueling Thompson sampling), the upper bound of the minimum width γ % credible interval of $f(x_2) - f(x_1)$ (dueling UCB) or a combination of the expected probability of improvement and the information gain (EIIG): $\kappa \log (\mu_c(x_1, x_2)) - I(c, f|\mathcal{D}_t, x_1, x_2)$. In this last acquisition function, κ determines a tradeoff between exploration and exploitation (here,

Name	Criterion
Binary Expected Improvement	$w_{}(x) = \mathbb{E}\left[\left[\Phi(f(x)) - \max(u(x))\right]\right]$
(Tesch et al., 2013)	$u_{\text{BEI}}(\mathbf{x}) = \mathbb{E}\left[\left[\Psi(f(\mathbf{x})) - \max_{\mathbf{x} \in \mathcal{X}} (\mu_c(\mathbf{x}))\right]_+\right]$
Upper Credible Bound (latent space)	$\alpha_{\text{LICD}}(\mathbf{x}) = \mu_{c}(\mathbf{x}) + \beta_{c}\sigma_{c}(\mathbf{x})$
(Tesch et al., 2013)	$\mu_{\rm UCB_f}(\mathbf{x}) = \mu_f(\mathbf{x}) + \rho_f \sigma_f(\mathbf{x})$
Upper Credible Bound (probability space)	$\alpha_{\mathrm{UCB}_{\Phi}}(\mathbf{x}) = \mathbb{E}[\Phi(f(\mathbf{x})] + \beta \sqrt{\mathbb{V}[\Phi(f(\mathbf{x}))]}$
Binary Knowledge Gradient	$lpha_{\mathrm{KG}}(\mathbf{x}) = \mathbb{E}(\mu_{t+1}^{\star} - \mu_t^{\star} \mathscr{D}_t, \mathbf{x}_{t+1} = \mathbf{x})$

Table 3.1: *List of acquisition rules for binary Bayesian optimization. In blue are the rules introduced in this paper.*

we used $\kappa = 0.5$).

I will briefly cover existing acquisition rules (summarized in table 3.2), approximately following the chronological order of their publication. Based on the analysis of their advantages and drawbacks, we will also propose new acquisition rules and compare their performance in synthetic experiments.

Improvement-based policies

Bivariate Expected Improvement For preferential Bayesian optimization, Brochu et al. (2010b) proposed to choose as the first duel member the maximum of the latent function mean among previously evaluated points: $x_1 = x^{\star\star} = \underset{x \in X_t}{\arg \max} \mu_g(x)$, and as the second duel member the maximum of the latent variance: $x_2 = \underset{x \in X}{\arg \max} \sigma_g(x)$. The idea is to select a challenger whose value is the most uncertain. However, as discussed in 3.3.2, high latent variance $\sigma_g^2(x)$ in conjunction with high latent mean $\mu_g(x)$ can correspond to a little uncertain output. To see this, consider the case where x_2 is such that $\mu_g(x_2) << \mu_g(x_1)$. Since the expected output of the duel is $\mu_c(x_1, x_2) = \Phi\left(\frac{\mu_g(x_1) - \mu_g(x_2)}{\sqrt{1 + \sigma_f^2(x_1, x_2)}}\right)$, even if $x_2 = \underset{x \in \mathcal{X}}{\arg \max} \sigma_g(x)$, we can have $\mu_c(x_1, x_2) \approx 1$ (see figure 3.10). So the result of the duel will most likely be $x_1 \succ x_2$. Since the output of the duel is little uncertain, this query is not explorative.

Brochu et al. (2008, 2010b) then proposed to generalize expected improvement to PBO. Improvement was defined² as:

$$I(x) = \max(0, g(x) - \mu_g^{\star\star}),$$
(3.37)

to which would correspond the expected improvement as computed in (Brochu et al., 2008, 2010a):

$$\alpha_{\rm EI}(\mathbf{x}) = \begin{cases} \mu_I(\mathbf{x}) \Phi\left(\frac{\mu_I(\mathbf{x})}{\sigma_g(\mathbf{x})}\right) + \sigma_g(\mathbf{x}) \phi\left(\frac{\mu_I(\mathbf{x})}{\sigma_g(\mathbf{x})}\right) & \text{if } \sigma_g(\mathbf{x}) > 0\\ 0 & \text{if } \sigma_g(\mathbf{x}) = 0 \end{cases}$$
(3.38)

with $\mu_I(\mathbf{x}) = \mu_g(\mathbf{x}) - \mu_g^{\star\star}$.

However, as noted by Nielsen et al. (2014) and Lorenz (2017), this acquisition function can lead to the undesirable effect of querying the duel ($x^{\star\star}, x^{\star\star}$), which does not bring any information and causes the algorithm to query the same duel over and over. This stems from the fact that this definition of improvement neglects the covariance between x_1 and x_2 . This behavior is illustrated in figure 3.16.

To solve this problem, Nielsen et al. (2014) proposed to use expected improvement for PBO, but keeping the original definition of improvement:

$$I(x) = \max(0, g(x) - g(x^{**}))$$
(3.39)

This definition takes explicitly into account the uncertainty about $g(x^{\star\star})$. The corresponding acquisition criterion is:

$$\alpha_{\mathrm{EI}}(\mathbf{x}) = \mathbb{E}(I(\mathbf{x})) = \int_{I=0}^{I=\infty} I \mathcal{N}(\mu_I(\mathbf{x}), \sigma_I^2(\mathbf{x})) dI, \qquad (3.40)$$

²The original definition is $I(\mathbf{x}) = \max(0, \mu_g(\mathbf{x}) - \mu_g^{\star\star})$. We believe this is a typo repeated in the two papers. Indeed, when using this definition, since μ_g and $\mu_g^{\star\star}$ are both deterministic in this context, we would get $\alpha_{\text{EI}}(\mathbf{x}) = \max(0, \mu_g(\mathbf{x}) - \mu_g^{\star\star})$, which does not correspond to the authors' results in equation 3.38.



Figure 3.16: Pathological behavior of the original Expected Improvement acquisition function by Brochu et al. (2008, 2010a). A. Queries of the acquisition function throughout the optimization experiment. In blue is indicated the first duel member, and in blue is the second duel member. After about 15 iterations, Expected Improvement selects identical duel members. As a consequence, the output is 1 with a probability of 0.5, and the new observation will not bring new information, so the acquisition function will keep selecting the same duel over and over. B. Posterior distribution of the Gaussian process surrogate model. Shaded areas correspond to ventiles of the distribution. The maximum of the latent mean is indicated by the blue dotted line. C. In this example, the Expected Improvement acquisition function (in blue) selects identical duel members, corresponding to the maximum of the posterior mean. On the contrary, Bivariate Expected Improvement (in red) selects as first duel member the maximum of the posterior mean, and as second duel member, the point indicated with the red dotted line. The difference comes from the fact that Bivariate Expected Improvement considers the covariance between the estimated maximum and the new candidate.

where $\sigma_I^2(\mathbf{x}) = \Sigma(\mathbf{x}, \mathbf{x}) + \Sigma(\mathbf{x}^{\star\star}, \mathbf{x}^{\star\star}) - 2\Sigma(\mathbf{x}, \mathbf{x}^{\star\star})$ and $\mu_I(\mathbf{x}) = \mu_g(\mathbf{x}) - \mu_g^{\star\star}$. The integral in 3.40 can be analytically computed:

$$\alpha_{\rm EI}(\mathbf{x}) = \mu_I(\mathbf{x})\Phi\left(\frac{\mu_I(\mathbf{x})}{\sigma_I(\mathbf{x})}\right) + \sigma_I(\mathbf{x})\phi\left(\frac{\mu_I(\mathbf{x})}{\sigma_I(\mathbf{x})}\right)$$
(3.41)

Nielsen et al. (2014) termed this acquisition rule "Bivariate Expected Improvement". An advantage compared to the previous version of expected improvement is that here, the covariance term $\Sigma(x, x^{\star\star})$ prevents from querying the duel ($x^{\star\star}, x^{\star\star}$).

Copeland Expected Improvement Later, Gonzalez et al. (2017) proposed a new acquisition function, based on the concept of the Copeland score. When the Copeland score is used for the identification step, a natural acquisition rule is thus to maximize the one-step ahead expected increase in the Copeland score of the Condorcet winner. The "Copeland Expected Improvement" is defined as:

$$\alpha_{\text{CEI}}(\mathbf{x}_{1}, \mathbf{x}_{2} \mid \mathscr{D}_{t}) = \mathbb{E}_{p(c|\mathbf{x}_{1}, \mathbf{x}_{2}, \mathscr{D}_{t})} \left[\left[C(\mathbf{x}_{c,t+1}) - C(\mathbf{x}_{c,t}) \right]_{+} \mid \mathscr{D}_{t} \right] \\ = \mu_{c}(\mathbf{x}, \mathbf{x}') \left[c_{0,t+1} - C(\mathbf{x}_{c,t}) \right]_{+} + \left(1 - \mu_{c}(\mathbf{x}, \mathbf{x}') \right) \left[c_{1,t+1} - C(\mathbf{x}_{c,t}) \right]_{+},$$
(3.42)

where $c_{1,t+1}$ (resp $c_{0,t+1}$) is the Copeland score of the new Condorcet winner if the output of the new duel is $x \succ x'$ (resp. $x' \succ x$). However, α_{CEI} is expensive to compute as it requires estimating the Condorcet winner for each possible output of the duel, which itself has a cost

that grows exponentially with the dimension of the search space. Gonzalez et al. (2017) used a discretization of the search space to estimate its maximum. Its use is thus limited to lowdimensional cases.

Here, we suggest to make direct use of the posterior distribution over *g*, by using a preferential version of the knowledge gradient, which greedily minimizes the expected one-step ahead regret when the identification step is $\arg \max_{x \in \mathcal{X}} \mu_g(x)$:

$$\alpha_{\mathrm{PKG}}(\mathbf{x}_1, \mathbf{x}_2) = \mathbb{E}_{p(c|\mathbf{x}_1, \mathbf{x}_2, \mathscr{D}_t)} \left[\max_{\mathbf{x} \in \mathcal{X}} \mu_g(\mathbf{x}|\mathcal{D}_t, \mathbf{x}_1, \mathbf{x}_2) - \max_{\mathbf{x} \in \mathcal{X}} \mu_g(\mathbf{x}|\mathcal{D}_t) \right]$$
(3.43)

Note that this is identical to the Copeland Expected Improvement, but the soft-Copeland score of the Condorcet winner is replaced with the maximum posterior mean of the latent GP. Computing this maximum is much easier for high dimensions than computing the Condorcet winner. Here, the expectation is taken over the output of the new duel. So:

$$\alpha_{\mathrm{PKG}}(\mathbf{x}_1, \mathbf{x}_2) = \mu_c(\mathbf{x}_1, \mathbf{x}_2) \max_{\mathbf{x} \in \mathcal{X}} \mu_g(\mathbf{x} | \mathscr{D} \cup (\mathbf{x}_1 \succ \mathbf{x}_2)) + \mu_c(\mathbf{x}_2, \mathbf{x}_1) \max_{\mathbf{x} \in \mathcal{X}} \mu_g(\mathbf{x} | \mathscr{D} \cup (\mathbf{x}_2 \succ \mathbf{x}_1))$$
(3.44)

Contrary to CEI, this acquisition rule has a tractable gradient when using Laplace approximation (see appendix E). However, even with these modifications, it is impractically slow as each function evaluation requires solving two optimization problems (one for each possible output of the duel). As a consequence, we did not consider these joint acquisition methods on benchmarks (section 3.4).

Heuristics based on Thompson sampling

Duel Thompson Sampling Gonzalez et al. (2017) also proposed an efficient acquisition function, denoted "Duel Thompson Sampling". The first duel member x_1 is selected in an exploitative step using Thompson sampling, i.e., it is a sample from $p(\arg \max C(x)|\mathscr{D}_t)$. The second duel member, selected in an exploratory step, is $\arg \max \mathbb{V}[\Phi(f(x_1, x_2))]$. Thompson sampling is performed by sampling a linear approximation to the preference function \tilde{f} from the Gaussian process (see section 3.6 for details), and computing the Condorcet winner: $x_1 = \arg \max_{x \in \mathcal{X}} \int_{\mathcal{X}} \Phi(\tilde{f}(x_1, x_2)) dx'$. However, since the computation of the soft-Copeland score requires to integrate $\Phi(\tilde{f}(x_1, x_2))$ over the whole search space \mathcal{X} , the use of this acquisition function is restricted to low-dimensional problems.

Note that, in the paper, the method proposed to sample from $p(f|\mathcal{D}_t)$ does not preserve the anti-symmetric property of the preference function (see figure 4 in Gonzalez et al. (2017)), leading to samples that do not satisfy $\tilde{f}(x_1, x_2) = -\tilde{f}(x', x)$. We will discuss this issue in section 3.6.2 and explain how to address it.

Moreover, note that the first step is equivalent to sampling from $p(\mathbf{x}^*|\mathcal{D}_t)$, or equivalently maximizing a sample \tilde{g} from the posterior Gaussian process since we have the following property (see the proof in appendix H):

$$\underset{\mathbf{x}\in\mathcal{X}}{\arg\max} \widetilde{g}(\mathbf{x}) = \underset{\mathbf{x}\in\mathcal{X}}{\arg\max} \int_{\mathcal{X}} \Phi(\widetilde{f}(\mathbf{x}_1, \mathbf{x}_2)) d\mathbf{x}'$$

So we suggest to compute the same acquisition function at lower cost by replacing the

selection of x_1 with: $x_1 = \underset{x \in \mathcal{X}}{\operatorname{arg\,max}} \widetilde{g}(x)$. This modification allows to use Duel Thompson Sampling in higher dimensional problems.

Here, when comparing acquisition functions, we used this modified yet equivalent version of Duel Thompson Sampling. As we will see in section 3.6, Thompson sampling with GP classification models comes with several challenges, and the original method proposed by Gonzalez et al. (2017) has several unwanted properties. We will discuss methods to efficiently solve these issues.

KernelSelfSparring Sui et al. (2017) proposed KernelSelfSparring, a generalisation of Thompson sampling in which inputs are sampled from $p(x^*|\mathscr{D}_t)$. For example, in the case of duels, two samples \tilde{g}_1 and \tilde{g}_2 are drawn from the GP posterior $p(g|\mathscr{D}_t)$ and their maximum $x_1 = \underset{x \in \mathcal{X}}{\arg \max \tilde{g}_1(x)}$ and $x_2 = \underset{x \in \mathcal{X}}{\arg \max \tilde{g}_2(x)}$ are taken as members of the next duel. This acquisition function was initially designed for applications to multi-dueling bandits with correlated arms, that is, multi-armed bandits where feedbacks come in the form of comparisons between groups of inputs. In this paper, however, the Gaussian process model directly modelled the probabilities of preference judgements (not a latent function). A drawback of KernelSelfSparring is that it selects the duel members independently from each other. As a consequence, it can select two highly correlated duelists in duels that are not very informative. For example, it can select identical duel members.

Information-based policies

Thatte et al. (2017) applied the Predictive Entropy Search acquisition rule (Hernández-Lobato et al., 2014) to PBO. The Predictive Entropy Search with Preferences (PES-P) selects comparisons that optimally reduce the uncertainty in the distribution of the objective function maximum with the least number of queries, or equivalently, that maximizes the mutual information between the output of a duel and the maximum x^* of g. That is, it selects:

$$\underset{(\boldsymbol{x}_1,\boldsymbol{x}_2)\in\mathcal{X}^2}{\arg\max} I(c,\boldsymbol{x}^*|\mathscr{D}_t,\boldsymbol{x}_1,\boldsymbol{x}_2), \tag{3.45}$$

which can be computed by noting that:

$$I(c, \mathbf{x}^{\star} | \mathscr{D}_t, \mathbf{x}_1, \mathbf{x}_2) = H(c | \mathscr{D}_t, \mathbf{x}_1, \mathbf{x}_2) - \mathbb{E}_{p(\mathbf{x}^{\star} | \mathscr{D}_t)} H(c | \mathscr{D}, \mathbf{x}_1, \mathbf{x}_2).$$
(3.46)

This acquisition function has an intuitive and elegant interpretation: it selects the duel that is expected to reduce our uncertainty about x^* the most. However, this acquisition function is slow to compute as the second term in 3.46 requires to sample from $p(x^*|\mathscr{D})$ and condition the GP on x^* , which necessitates several approximations.

Maximally Uncertain Challenge

A problem with previously described acquisition functions is that they are either computationally expensive and do not scale well with dimensionality, or do not explicitly consider the tradeoff between exploration and exploitation. This is for example the case of KernelSelf-Sparring which, as mentioned earlier, can select two highly correlated duelists. Intuitively, following the rationale of improvement-based policies, one may want to select duels that are likely to change our belief about the current maximum. This can be performed by selecting the maximum of the posterior mean and the challenger for which the output of the duel is most epistemically uncertain.

I thus propose a new sequential acquisition rule, which is a modification of the Duel Thompson Sampling acquisition rule (Gonzalez et al., 2017), and that we call the Maximally Uncertain Challenge (MUC, see figure 3.17), that selects the first duel member in an exploitative step by taking the maximum of the posterior mean

$$\mathbf{x}_1 = \mathbf{x}^\star = \underset{\mathbf{x} \in \mathcal{X}}{\arg \max} \ \mu_g(\mathbf{x}), \tag{3.47}$$

and then selects the challenger in an exploratory move, by maximizing epistemic uncertainty. That is:

$$\mathbf{x}_{2} = \underset{\mathbf{x}_{2} \in \mathcal{X}}{\arg \max} \mathbb{V}[\Phi(f(\mathbf{x}_{1}, \mathbf{x}_{2})) | \mathcal{D}_{t}].$$
(3.48)

Alternatively, one may consider selecting the challenger by maximizing the epistemic uncertainty defined as mutual information between *c* and $f(x_1, x_2)$. In pilot studies, the two approaches performed similarly, so we decided to focus on the representation of uncertainty based on variance.

Comparing acquisition functions

To compare the behavior of the different acquisition rules from the challenge class, one can visualize the Pareto front of the multi-objective maximization problem (figure 3.18):

$$\underset{\boldsymbol{x}\in\mathcal{X}}{\arg\max} \left(\mathbb{E}[\Phi(f(\boldsymbol{x},\boldsymbol{x}_{1}))|\mathcal{D}], \sqrt{\mathbb{V}[\Phi(f(\boldsymbol{x},\boldsymbol{x}_{1}))|\mathcal{D}]} \right).$$
(3.49)

Expected improvement (Brochu et al., 2010a) as well as Dueling Thompson tend to be much more exploitative than the other rules.

A good acquisition function should lead the optimization algorithm to rapidly converge towards the maximum of the objective. So, after spending the evaluation budget, the value of the reported solution should be as close as possible to the true maximum, and distance should rapidly decrease throughout the experiment. This can be measured using the Area Under Curve, that is, $\sum_{i=1}^{T} f(\mathbf{x}_{t}^{\star})$ (or $\sum_{i=1}^{T} g(\mathbf{x}_{t}^{\star})$ in the case of PBO).

The use of synthetic functions whose analytical form is known to compare algorithms is now widespread (Eggensperger et al., 2015; Bergstra and Bengio, 2012). We evaluated the various acquisition functions on a set of 34 functions from a widely used virtual library for optimization experiments (Surjanovic and Bingham, 2021). The functions in this library are fast to compute, but exhibit a diversity of behaviors that occur in real-life optimization problems. To perform statistical analysis of acquisition functions performance on the benchmarks, we used the stratified analysis method proposed by Dewancker et al. (2016). Briefly, for each benchmark function, we performed pairwise comparisons between acquisition functions using the Mann-Whitney U test at $\alpha = 5 \times 10^{-4}$ significance on the best value found at the end of the optimization sequence. This determines a partial ranking based on the number of wins.

Class	Name	First member	Second member	
	Expected Improvement (Brochu et al., 2010b)	<i>x</i> **	$\underset{\boldsymbol{x}\in\mathcal{X}}{\arg\max} \mathbb{E}\left[\left[g(\boldsymbol{x}) - \mu_{g}^{\star\star}\right]_{+} \mathscr{D}_{t}\right]$	
Challenge	Bivariate Expected Improvement (Nielsen et al., 2014)	<i>x</i> **	$\underset{\boldsymbol{x}\in\mathcal{X}}{\arg\max} \mathbb{E}\left[\left[g(\boldsymbol{x}) - g\left(\boldsymbol{x}^{\star\star}\right)\right]_{+} \mathscr{D}_{t}\right]$	
	Maximally Uncertain Challenge (ours)	<i>x</i> *	$\underset{\mathbf{x}\in\mathcal{X}}{\arg\max} \mathbb{V}\left[\Phi(f(\mathbf{x}))\right]$	
	Dueling Thompson (Benavoli et al., 2020b)	x^{\star}	Sample from $p(\mathbf{x}^{\star} \mathscr{D}_{t})$	
	Dueling UCB (Benavoli et al., 2020b)	x^{\star}	$\arg\max_{\boldsymbol{x}\in\mathcal{X}} \left[\mu_g(\boldsymbol{x}) + \beta\sigma_g(\boldsymbol{x})\right]$	
	EI with Information Gain (Benavoli et al., 2020b)	x^{\star}	$\underset{x_2 \in \mathcal{X}}{\arg\max k \log(\mu_c(\mathbf{x}_2, \mathbf{x}^*))} - I(c, f \mathcal{D}_t, \mathbf{x}^*, \mathbf{x}_2)$	
Thompson sampling	Duel Thomspon Sampling Gonzalez et al. (2017)	$\boldsymbol{x} \sim p\left(\boldsymbol{x}^{\star} \boldsymbol{\mathcal{D}}_t\right)$	$\underset{\textbf{x}\in\mathcal{X}}{\arg\max} \mathbb{V}\left[\Phi(f(\textbf{x}))\right]$	
	Kernel Self Sparring Sui and Burdick (2017)	$\boldsymbol{x} \sim p\left(\boldsymbol{x}^{\star} \mathscr{D}_{t}\right)$	$oldsymbol{x} \sim p\left(oldsymbol{x}^{\star} \mathscr{D}_t ight)$	
Joint selection	Copeland Expected Improvement (Gonzalez et al., 2017)	$\arg\max_{(x_1,x_1)\in\mathcal{X}}$	$\sum_{2} \mathbb{E}_{c} \left[\left[C(\boldsymbol{x}_{c,n+1}) - C(\boldsymbol{x}_{c,n}) \right]_{+} \mid \mathcal{D}_{t}, \boldsymbol{x}_{1}, \boldsymbol{x}_{2} \right]$	
,	Preference Knowledge Gradient	$\arg\max_{(\mathbf{x}_1, \mathbf{x}_2) \in \mathcal{X}^2} \mathbb{E}_{\boldsymbol{\mathcal{X}}} \left[\max_{\mathbf{x} \in \mathcal{X}} \left(\mu_g \left(\mathbf{x} \mathscr{D}_t, \mathbf{x}_1, \mathbf{x}_2 \right) \right) \right]$		
F	Predictive Entropy Search with Preferences (Thatte et al., 2017)	$\arg \max_{(x_1,x_2)\in\mathcal{X}^2} I\left(x^*,c\big \mathscr{D}_t,x_1,x_2\right)$		

Table 3.2: List of acquisition rules for Preferential Bayesian Optimization. In blue are the rules introduced in this paper. Acquisition functions can be divided into three main classes. The first class corresponds to 'challenges', i.e. when the first duel member is the maximum of the posterior mean of the latent GP x^* , sometimes restricted to training points (noted x^{**}). In the second class, the first duel member is selected via Thompson sampling. For both classes, the second duel member is selected using various heuristics that attempt to tradeoff exploration and exploitation. In the third class, the two duelists are jointly selected.



Figure 3.17: Maximally Uncertain Challenge (MUC). A. Preference relationships. White circles correspond to observations $\mathbf{x} \succ \mathbf{x}'$, and black circles correspond to $\mathbf{x}' \succ \mathbf{x}$. B. Posterior latent GP (in blue) and true value function (in red). The GP in conditioned on $g(\mathbf{x}_0) = 0$. The MUC acquisition rule selects as first duel member \mathbf{x}_1 the maximum of the GP mean. C. Probability that a challenger wins the duel against \mathbf{x}_1 (in red), and posterior probability of \mathbf{x} winning against \mathbf{x}_1 . D. Epistemic uncertainty on the output of a duel against \mathbf{x}_1 , as measured by $\mathbb{V}[\Phi(f(\mathbf{x}, \mathbf{x}_1))]$ (in blue). Total uncertainty is drawn in red. The MUC acquisition rule selects a challenger \mathbf{x}_2 such that the output of the duel against the 'champion' \mathbf{x}_1 is most epistemically uncertain. Note that this does not correspond to the duel maximizing total uncertainty.

Ties are then broken by running the same procedure, but based on the Area Under Curve, which is related to the speed at which the algorithm reaches the optimum. This generates a new partial ranking, based on which a Borda score (Dwork et al., 2001) is attributed to each acquisition function (the Borda score of a candidate is the number of candidates with a lower rank). Then, rankings from different benchmarks are aggregated by summing the Borda scores to establish a global ranking. This can be seen as a weighted vote from each benchmark function. One of the advantages of this method is that it allows aggregating the results over a large set of benchmark functions.

3.4 Results

For each objective function in a set of 34 functions from a widely used virtual library for optimization experiments (Surjanovic and Bingham, 2021), we inferred the hyperparameters for



Figure 3.18: Behavior of acquisition functions of the 'challenge' class. A. Acquisition functions from the challenge class select the maximum of the posterior mean as the first duel member \mathbf{x}_1 (in some cases, the maximum is selected among the training points, but for this plot we took the maximum on the whole search space). The second duel member is selected according to different heuristics, and is indicated by markers. **B.** The exploration-exploitation tradeoff can be visualized in the $(\mathbb{E}[\Phi(f(\mathbf{x},\mathbf{x}_1))|\mathscr{D}], \mathbb{V}[\Phi(f(\mathbf{x},\mathbf{x}_1))|\mathscr{D}])$ plane. The Pareto front of the corresponding multi-objective maximization problem is indicated in red and reported on the curve $P(\mathbf{x} \succ \mathbf{x}_1)$ in A. Expected Improvement as defined by Brochu et al. (2010b) and Dueling Thompson sampling tend to underexplore. Although this is not visible here, dueling Thompson is not guaranteed to select points from the Pareto front. In this example, the maximally uncertain challenge and dueling UCB select the same challenger by maximizing epistemic uncertainty about the duel outcome.

three different kernels (squared exponential, Matérn 3/2 and Matérn 5/2, see appendix C for details on these kernels) using maximum likelihood estimation with 1000 randomly chosen samples. We then determined for each function the kernel that was better at describing the function by measuring the root-mean-squared error on 3000 points. The benchmark functions are listed in appendix J.

3.4.1 Binary Bayesian optimization

I ran optimizations on each test function for 60 different random number generator seeds, for 100 iterations. In all cases except the knowledge gradient acquisition functions, I used the Expectation Propagation approximation. Indeed, the gradient of the knowledge gradient seems intractable with this approximation, so I used the Laplace approximation in this case.

To avoid saturation effects when transforming the benchmark functions through the nonlinearity, I scaled the test functions so that their mean is 0 and variance 1. The results are presented in table 3.3. More detailed results showing pairwise comparisons between acquisition functions are presented in the supplementary 3.26. The evolution of the value of the inferred maxima for five test functions is plotted in figure 3.19.

The UCB_{Φ} that we introduced outperforms all the other acquisition functions. In particular, in agreement with the results obtained by Tesch et al. (2013) for expected improvement, UCB defined in the latent space (UCB_{*f*}) do not perform well compared to its couterpart defined in the probability space (UCB_{Φ}), and to Binary Expected Improvement. We did not investigate the influence of the value of β on the performance of UCB_{Φ}. However, the fact that the perfor-

Acquisition rule	Rank	Borda score
UCB_{Φ}	1	105
Binary Expected Improvement (Tesch et al., 2013)	2	65
UCB_f (Tesch et al., 2013)	3	62
Thompson sampling	4	54
Binary Knowledge Gradient	5	32
Random	6	20

Chapter 3. Practical Bayesian optimization with binary and preference data

Table 3.3: Comparison of acquisition functions on benchmarks in binary Bayesian optimisation. UCB_{Φ} outperforms the other acquisition functions.

mance was good with an arbitrarily chosen value suggests that it may be robust to change in β . Surprisingly the performance of the Binary Knowledge Gradient (BKG) (in terms of Borda score) is very closed to the one of random acquisition. One possible reason is that we defined BKG so as to explicitly improve the latent mean. Defining a version of BKG aiming at maximizing the maximum probability of success $\mu_c(x_t^*)$ may lead to better results, but we did not investigate this possibility.



Figure 3.19: Evolution of the value of the inferred maximum $g(\mathbf{x}^*)$ in five experiments of binary Bayesian optimization, repeated 60 times on 100 iterations. The ranking between acquisition functions varies depending on the objective function, however, statistical analysis (table 3.3) shows that UCB_{Φ} and EI_{Φ} find better estimates of the maximum more rapidly.

3.4.2 Preferential Bayesian optimization

The results are presented in table 3.4.

We ran similar experiments, repeated with 40 different random seeds for 80 iterations each. Again, more detailed results showing pairwise comparisons between acquisition functions are presented in the supplementary 3.9. The Maximally Uncertain Challenge acquisition rule outperforms all the other rules. The dueling UCB rule shows very similar performance, which is surprising since it is defined in the latent space. This suggests that a prefential version of UCB_{Φ} may lead to even better performance, but we did not investigate this possibility. Bivariate Expected Improvement displays good performance, in particular compared to Expected Improvement (EI). Indeed, because of the frequent occurrence of pathological behavior (figure 3.16), EI shows results only slightly better than random acquisition. In general, acquisition rules based on Thompson sampling (Duel TS, Dueling Thompson, KernelSelfSparring) have limited performance.

Acquisition rule	Rank	Borda score
Maximally Uncertain Challenge	1	141
Dueling UCB (Benavoli et al., 2020a)	1	141
Bivariate Expected Improvement (Nielsen, 2015)	3	111
Duel Thompson Sampling (modified from Gonzalez et al. (2017)	4	106
Dueling Thompson (Benavoli et al., 2020a)	5	101
KernelSelfSparring (Sui and Burdick, 2017)	6	66
EIIG (Benavoli et al., 2020a)	7	32
Expected Improvement (Brochu et al., 2010b)	8	23
Random	9	21

Table 3.4: Comparison of acquisition functions on benchmarks.

3.5 Generalization to batches

In many preference learning situations, the user is asked to compare *m* configurations, where m > 2. This is called multi-dueling (Sui and Burdick, 2017), or batch preference learning (Siivola et al., 2020). In practice, performing batch comparisons allows gathering more information at each query. For example, consider a movie recommendation system trying to learn the preference of individual users. The system presents the users with a batch of *m* movies and they pick their preferred one from the list. Assuming responses are not stochastic and without prior information, in a single query, the system learns the preferred movie from the list, whereas with pairwise comparisons this would have required m - 1 queries.

Different feedback scenarios are possible, between two extremes: either we perform the $\frac{m(m-1)}{2}$ pairwise comparisons, either the output is the preferred configuration among the *m* possibilities. Here, we consider the case where we collect comparisons between each pair of variables. In this case, the output of each comparison (i, j) can be treated as a binary variable, $c_{ij} \in \{0, 1\}$, sampled from a distribution:

$$P(c_{ij}|f(x_i), f(x_j)) = \Phi(f(x_i) - f(x_j))$$
(3.50)

105



Figure 3.20: Evolution of the value of the inferred maximum $g(\mathbf{x}^*)$ in five experiments of preferential Bayesian optimization, repeated 20 times on 50 iterations. The ranking between acquisition functions varies depending on the objective function, however, statistical analysis (table 3.4) shows that over the whole functions set, the Dueling UCB and Maximally Uncertain Challenge find better estimates of the maximum more rapidly.

Since, conditioned on *f*, observations are independent:

$$P(\mathbf{c}|f(\mathbf{x}_1),\ldots,f(\mathbf{x}_m)) = \prod_{i=1}^{m-1} \prod_{j>i} P(c_{ij}|f(\mathbf{x}_i),f(\mathbf{x}_j)),$$
(3.51)

where *c* is a vector containing the output of all pairwise comparisons. In this case, the problem reduces to a standard preference learning problem with m(m-1)/2 pairwise comparisons, and can be treated using the same inference methods as standard GP preference learning Note however, that this decomposition does not hold for feedback scenarios involving ranking between three or more inputs, such as $x_1 \succ x_2 \succ x_3$ (Nguyen et al., 2020).

In the Gaussian process bandits setting, Sui and Burdick (2017) proposed the use of Kernel-SelfSparring, which is a variant of Thompson sampling that consists in sampling *m* candidates from $p(x^*|\mathscr{D}_t)$ and make them compete. Siivola et al. (2020) independently proposed the same acquisition rule and showed that it outperformed the generalization of Expected Improvement to batches.

An obvious drawback of this acquisition rule is that it selects points independently from each other. In particular, it does not prevent sampling duels between points that are highly correlated, leading to under-exploratory behavior. A better acquisition function would try to select points that have little correlation. Given the much better performance of the Maximally Uncertain Challenge acquisition rule compared to KernelSelfSparring in the case of m = 2, we



Figure 3.21: Batch preferential Bayesian optimization. *A.* With batches of size 4. Note that the Maximally Uncertain Challenges acquisition rule is more explorative than KernelSelfSparring. *B.* With batches of size 25. For large batches, the generalization of MUC tend to select the same challengers several times.

generalized it to the batch case. The first competitor is $x_1 = x^* = \underset{x \in \mathcal{X}}{\operatorname{arg max}} \mu_g(x)$. A batch of challengers is then selected in an exploratory move, as:

$$(\mathbf{x}_{2},\cdots,\mathbf{x}_{m}) = \underset{\mathbf{x}_{2},\cdots,\mathbf{x}_{m}\in\mathcal{X}}{\arg\max} \sum_{i=1}^{m} \sum_{j>i} \mathbb{V}[\Phi(f(\mathbf{x}_{i},\mathbf{x}_{j}))]$$
(3.52)

This way, we select challengers for which we are the most uncertain about the outcome of the duel with x_1 and the duels between the challengers (figure 3.21).

To compare this new acquisition rule to KernelSelfSparring, we ran the optimization with m = 3 for 30 iterations and 30 repetitions. We ran the experiment for the feedback scenario where all pairwise comparisons are observed. The results are presented on table 3.5. The generalization of the Maximally Uncertain Challenge rule outperforms KernelSelfSparring. This is likely because KernelSelfSparring selects batch members independently from each other, causing redundancy.

However, the batch variant of MUC does not consider the correlation between the corresponding selected duels. As a consequence, for large batches, this acquisition function tends to select the same challenges several times.

Acquisition rule	Rank	Borda score
Maximally Uncertain Challenges	1	31
KernelSelfSparring (Sui and Burdick, 2017)	2	22
Random	3	0

Table 3.5: Comparison of acquisition functions on benchmarks in batch preferential optimization. The batch size was 3. We ran the experiment with the type of feedback where all pairwise comparisons between batch members are observed. The MUC acquisition rule largely outperforms KernelSelfSparring.

The problem of batch optimization is tightly linked to non-myopic optimization, where the acquisition step takes into account the fact that future queries will follow. Recently, much research efforts have been devoted to the problem of non-myopic BO (Wu and Frazier, 2019; Jiang et al., 2019), as it can dramatically increase sample-efficiency in some situations. The link between batch and non-myopic optimization comes from the fact the optimal batch of size *m* is an approximation to the *m*-steps look-ahead optimal points (Ginsbourger et al., 2008; Azimi et al., 2012; Jiang et al., 2019, 2020). We believe this new acquisition method could be efficient for non-myopic preferential optimization. However, we did not investigate this problem, which should be the topic of future work.

3.6 Efficiently sampling from the maximum posterior distribution

In the following section, we will present some practical considerations regarding the approximate sampling from Gaussian process posteriors with non-Gaussian likelihood, which is involved in several acquisition rules as well as in other applications such as GP-based simulators (Wilson et al., 2020a), which are used for example in model-based policy search in robotics (Deisenroth et al., 2015).

3.6.1 Efficiently sampling from Gaussian processes

Acquisition rules based on Thompson sampling (Thompson, 1933; Russo et al., 2018): Kernel-SelfSparring (Sui et al., 2017), Dueling Thompson Sampling (Benavoli et al., 2020b), and Duel Thompson Sampling (Gonzalez et al., 2017), as well as Predictive Entropy Search (Hernández-Lobato et al., 2014) rely on samples from the posterior distribution over the maximum (see table 3.2):

$$p(\mathbf{x}^{\star} \mid \mathscr{D}_{t}) = p\left(f(\mathbf{x}^{\star}) = \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \mid \mathscr{D}_{t}\right)$$
(3.53)

Hernández-Lobato et al. (2014) proposed the following sampling scheme: draw a sample from the posterior distribution $p(f|\mathcal{D}_t)$, then return the maximum of the sample. One could iteratively construct the sample f while it is being optimized but, as noted by Hernández-Lobato et al. (2014), this would have a cost $\mathcal{O}(m^3)$, where m is the number of evaluations of the function necessary to find the maximum. Although this is doable in practice (see e.g. Siivola et al. (2020)), Hernández-Lobato et al. (2014) suggested a more efficient procedure by sampling a finite-dimensional approximation to f, based on a finite-dimensional approximation to the kernel $k(x, x') \sim \phi(x)^{\top} \phi(x')$ (Lázaro-Gredilla et al., 2010). In GP classification and preference learning, this approximate sampling cannot be directly applied. In the following section, we will explain how to apply existing approximate sampling methods to the case of GP classification models.

Weight-space approximate sampling

The most widely used method for approximate sampling from GP in Bayesian optimization is the weight-space approximation. Assume that we have a finite-dimensional approximation to the kernel $k(x, x') \sim \phi(x)^{\top} \phi(x')$ (I describe methods for approximating kernels in 3.6.2). The features $\phi(x)$ can be used to approximate the Gaussian process posterior with a Bayesian linear model: $f(x) \sim \phi(x)^{\top} \omega$, where (Lázaro-Gredilla et al., 2010):

$$\omega \sim \mathcal{N}\left(\left(\Phi^{\top}\Phi + \sigma^{2}I\right)^{-1}\Phi^{\top}\boldsymbol{y}, \left(\Phi^{\top}\Phi + \sigma^{2}I\right)^{-1}\sigma^{2}\right)$$
(3.54)

See details and the proof in appendix G).

However, the degeneracy, i.e., low-rankness of the GP approximation, causes the estimate to grow over-confident when the number of observed points exceeds the degrees of freedom of the approximation. This results in ill-behaved approximations, and, in particular, underestimated variance, in regions far away from the data points. This phenomenon is known as variance starvation (Wang et al., 2017a; Mutný and Krause, 2018; Calandriello et al., 2019).

Decoupled-bases approximate sampling

Recently Wilson et al. (2020a) proposed a more efficient way to sample from Gaussian process posteriors that avoids variance starvation. The original sampling method was devised for exact GP with Gaussian noise and sparse Gaussian processes, where the GP is computed based on a set of inducing points that explain the data, however, it can easily be generalized to non-Gaussian likelihood with a latent function (Wilson et al., 2020b).

Briefly, this method is based on a corollary of Matheron's rule. For a Gaussian process $f \sim \mathcal{GP}(0, k)$, the latent process conditioned on latent values (X, y) admits, in distribution, the representation:

$$\underbrace{\left(f \mid \boldsymbol{y}\right)\left(\cdot\right)}_{\text{posterior}} \stackrel{\text{d}}{=} \underbrace{f\left(\cdot\right)}_{\text{prior}} + \underbrace{k\left(\cdot, \boldsymbol{x}\right)\mathbf{K}^{-1}\left(\boldsymbol{y} - f(\boldsymbol{X})\right)}_{\text{update}}$$
(3.55)

This corollary defines an approximation to the Gaussian process conditioned on (X, y), where the stationary prior is approximated with a Bayesian linear model (weight-space prior), and the approximate posterior is obtained by adding an exact update (function-space update):

$$\underbrace{(f \mid \boldsymbol{y})(\cdot)}_{\text{posterior}} \stackrel{\text{d}}{\approx} \underbrace{\sum_{i=1}^{\ell} \omega_i \phi_i(\cdot)}_{\text{weight-space prior}} + \underbrace{\sum_{j=1}^{m} v_j k\left(\cdot, \boldsymbol{x}_j\right)}_{\text{function space update}}, \qquad (3.56)$$

where $v = \mathbf{K}^{-1}(y - \Phi \omega)$, and ω is sampled from $\mathcal{N}(\mathbf{0}, I)$. This method is termed decoupled-bases decomposition of the Gaussian process (see figure 3.22).

To sample from the posterior latent function, we thus sample $\boldsymbol{\omega}$ from $\mathcal{N}(\mathbf{0}, I)$ and compute the corresponding weight-space prior, then sample \boldsymbol{y} from $\mathcal{N}(\mu_f(X), \Sigma_f(X, X))$ and compute the corresponding function-space update.

We will compare these two approximate sampling procedure in section 3.6.4. Now that we have considered procedures to sample finite-rank approximations to Gaussian processes, we need to define a way to approximate the kernel.

3.6.2 Kernel approximation

Sparse Spectrum Gaussian Process approximation

The most widely used method for kernel approximation in the GP literature is based on the fact that any shift-invariant kernel *k* can be approximated by means of the inner product of random features ϕ such that: $k(\mathbf{x}, \mathbf{x}') \sim \phi(\mathbf{x})^{\top} \phi(\mathbf{x}')$. This method is known as the random Fourier features algorithm (or random kitchen sinks) (Rahimi and Recht, 2008), and was applied to GP by Lázaro-Gredilla et al. (2010) under the name Sparse Spectrum GP approximation (SSGP).



Figure 3.22: Approximate GP sampling with the decoupled sampling method of Wilson et al. (2020a). The kernel is approximated with the Hilbert space method of Solin and Särkkä (2020). A. Red: true function. Red circles: data points. Blue: posterior mean of the GP, the shaded regions correspond to standard deviation. B. Pathwise update. An approximate sample is drawn from the GP prior (dark blue). The update (yellow) is computed exactly in the function space. The sample from the posterior GP (light blue) is obtained by summing the sample from the prior and the update, following the corollary of Matheron's rule. C. Approximate samples from the posterior GP. D. Distribution of the samples. Note the similarity with A.

The details of this approximation and a proof can be found for example in Hernández-Lobato et al. (2014), and are reproduced in appendix G.1.

Hilbert space Gaussian Process approximation

Recently, another method was proposed by Solin and Särkkä (2020), which aims at making the approximation as good as possible for a given rank (see Riutort-Mayol et al. (2020) for details about the practical implementation). As a consequence, for the same number of features, this reduced-rank GP approximation (RRGP) is better than the SSGP approximation.

In this method, the kernel is approximated using a series expansion in terms of eigenfunctions of the Laplace operator on a rectangular domain $\Omega = [-L_1, L_1] \times \cdots \times [-L_d, L_d]$ (the search space are usually rectangular in Bayesian optimization). See appendix G.2 for details.

3.6.3 Sampling from GP classification models

With non-Gaussian likelihoods, some aspects of the methods presented above need to be modified to account for the fact that the function values are not directly measured.

Weight-space approximation with non-Gaussian likelihoods

Gonzalez et al. (2017) applied the weight-space approximate sampling method to sample from GP posteriors in a preference model. However, in the case of non-Gaussian likelihoods, , naively replacing y in 3.54 by the latent values inferred by the Laplace approximation (see I) or Expectation Propagation would not take into account the correlated heteroscedastic noise on the latent function values at training points. To the extent of our knowledge, the process



Figure 3.23: Reduced rank approximations of the Matérn 5/2 kernel using the Hilbert space method of Solin and Särkkä (2020). The blue line corresponds to the true covariance with the Matérn 5/2 kernel, and the red lines correspond to approximations with 8, 16, 32 or 64 basis functions.

Algorithm 2: Approximate sampling from p (x* | D_t) (in blue: precomputed)
Input: A Gaussian process prior on g: GP (m (·), k (·, ·)), the corresponding posterior on f: f ~ GP (μ, Σ), D_t = (X, y), m
Compute the eigenfunctions φ_{j1},...,j_d and eigenvalues λ_{j1},...,j_d of the Laplace operator;
Compute the corresponding features for the preference kernel k_{pref} approximation;
Sample w ~ N(0, I) and u ~ N (μ (x, x₀), Σ ((x, x₀), (x, x₀)));
Compute the weight space prior w^Tφ_{pref} (x, x₀) and the function space update K⁻¹_{pref} (u - Φ_{pref}w) k_{pref} ((x, x₀), X);
Compute the corresponding approximate sample g̃ from the GP posterior and its gradient ;
Find the maximum x̃* of g̃ using algorithms such as L-BFGS ;
Output: A sample x̃* from p (x*|D_t)

of weight-space approximate sampling has not been rigorously introduced for latent Gaussian process models. Here, we suggest to use a sampling process in two steps. First, samples y are drawn from the posterior distribution over the latent variables at training points: $\mathcal{N}(\mu, \Sigma)$, then ω is sampled from $\omega \sim \mathcal{N}\left(\left(\Phi^{\top}\Phi + \sigma^{2}I\right)^{-1}\Phi^{\top}y, \left(\Phi^{\top}\Phi + \sigma^{2}I\right)^{-1}\sigma^{2}\right)$, where σ is a small constant used for regularization.

To see why this sampling scheme is correct, note that: $p(f(\mathbf{x})|\mathscr{D}_t) = \int p(f(\mathbf{x})|X, \mathbf{y})p(\mathbf{y}|\mathscr{D}_t)d\mathbf{y}$. So given latent values sampled from $p(\mathbf{y}|\mathscr{D}_t)$, an approximate sample \tilde{f} can be drawn from $p(f(\mathbf{x})|X, \mathbf{y})$ using the method of Hernández-Lobato et al. (2014).

Approximation of preference kernels

The sampling methods mentioned above consists in approximating a stationary kernel *k* by means of the inner product of features ϕ such that: $k(\mathbf{x}, \mathbf{x'}) \sim \phi(\mathbf{x})^{\top} \phi(\mathbf{x'})$. The base kernel used to model the value function may be shift-invariant, the preference kernel, however, is not in general (see the proof in appendix F.1).

This inexact hypothesis introduced in the sampling algorithm, although unnoticed by Gonzalez et al. (2017) who initially used it, leads to samples that are not consistent with the

anti-symmetric property of a preference function, i.e. f(x, x') = -f(x', x), (see figure 4.1 in the aforementioned paper).

However, assume that we have a finite dimensional approximation to the base kernel $k(\mathbf{x}, \mathbf{x}') \sim \phi(\mathbf{x})^{\top} \phi(\mathbf{x}')$, it is easy to see that we can approximate the preference kernel by $k_{\text{pref}}((\mathbf{x}_i, \mathbf{x}_j), (\mathbf{x}_k, \mathbf{x}_l)) \sim \phi_{\text{pref}}(\mathbf{x}_i, \mathbf{x}_j)^{\top} \phi_{\text{pref}}(\mathbf{x}_k, \mathbf{x}_l)$, with:

$$\phi_{\text{pref}}(\boldsymbol{x}_i, \boldsymbol{x}_j) = \phi(\boldsymbol{x}_i) - \phi(\boldsymbol{x}_j)$$
(3.57)

By construction, the corresponding sample is anti-symmetric.

The conditional preference kernel follows the same decomposition. However, the conditional base kernel \hat{k} is not stationary due to the conditioning. To circumvent this problem and sample from a weight-space prior with the appropriate conditioning, we use the following modification to the decoupled bases sampling scheme: first, (see 3 for the corresponding algorithm).

Algorithm 3: Approximate sampling from $p(x^*|\mathscr{D}_t)$ with a conditional preference kernel (in blue: precomputed)

Input: A Gaussian process prior on *g*: $\mathcal{GP}(0, \hat{k}(\cdot, \cdot))$, the corresponding posterior on *f*: $f \sim \mathcal{GP}(\mu, \Sigma), \mathcal{D}_t = (X, \mu), m$

Compute the eigenfunctions $\phi_{j_1,...,j_d}$ and eigenvalues $\lambda_{j_1,...,j_d}$ of the Laplace operator; Compute the corresponding features for the preference kernel k_{pref} approximation; Sample $\mathbf{w} \sim \mathcal{N}(0, I)$;

Compute the weight space prior $w^{\top}\phi(x)$ and the function space update

 $k(x_0, x_0)^{-1}(y_0 - w^{\top}\phi(x_0))k(x_0, x)$ corresponding to the observation (x_0, y_0) 3.56);

Compute the corresponding approximate sample \hat{g} and its gradient. \hat{g} is an

approximate sample from the prior GP conditioned on $g(x_0) = y_0$ on g. To \hat{g} corresponds \hat{f} ;

Sample $u \sim \mathcal{N}(\mu(\mathbf{x}, \mathbf{x}_0), \Sigma((\mathbf{x}, \mathbf{x}_0), (\mathbf{x}, \mathbf{x}_0)))$ and compute the corresponding function space update $\hat{K}_{\text{pref}}^{-1}(\mathbf{u} - \hat{f}(X))\hat{k}_{\text{pref}}((\mathbf{x}, \mathbf{x}_0), X)$ (eq. 3.56);

Compute the corresponding approximate sample \tilde{g} and its gradient ;

Find the maximum \tilde{x}^* of \tilde{g} using algorithms such as L-BFGS ;

Output: A sample \tilde{x}^* from $p(x^*|\mathcal{D}_t)$

3.6.4 Comparing approximate sampling methods

To compare the performance of these different sampling techniques in the case of GP classification models, following Wilson et al. (2020a), we measured the 2-Wasserstein distance between the true latent GP and the approximate samples distribution, estimated by drawing 5000 approximate samples. We repeated the measurement for 32 random functions drawn from Gaussian processes defined on [0, 1] with squared exponential kernel and for different sizes n of random training data. Given that we consider two methods for finite-rank approximation of kernels (RRGP and SSGP), and two methods to represent the approximation (the decoupled3.7. Software



Figure 3.24: Sampling from a preference GP A. Decoupled bases sampling from the value function with pathwise update. A approximate sample is drawn from the prior distribution (blue). An exact update (red) is computed and the sample is updated using Matheron's rule (turquoise), giving a sample from the posterior $p(g|\mathcal{D})$. B. Posterior mean of the GP model of the latent preference function. C. Posterior variance of the GP model of the latent preference function.D. Sample from the preference GP corresponding to the value function in A. Note that our sampling method guarantees that the sample is anti-symmetric. E. Average over 1000 samples from the preference GP, not the agreement with B. F. Variance over 1000 samples from the preference GP, note the agreement with C.

based approximation and the weight-space approximation), we have four possible approximation schemes. The number of features used to sample from the prior was m = 32. For fair comparison between the weight-space and the decoupled-bases methods, the number of features was n + m when using weight-space approximation. The results are plotted on figure 3.25. The reduced-rank kernel approximation is more accurate compared to sparse-spectrum approximation. Moreover, the decoupled-based approximation is more accurate than the weight-space approximation. These results for GP binary classification are in agreement with those obtained by Wilson et al. (2020a) and Solin and Särkkä (2020) for GP regression. In this work, we used the decoupled-based approximate sampling method in combination with the Hilbert-space kernel approximation method, as it is the most efficient approximate sampling method.

3.7 Software

Several open-source libraries for Bayesian optimization are available. However, only a fraction of them provides direct support for a scenario with binary outputs. Due to the popularity of deep learning frameworks, with features such as GPU acceleration and automatic differentiation, recent packages for Bayesian optimization are built on these libraries. BoTorch (Balandat et al., 2019a,b) is a Python library for Bayesian optimization based on GPytorch (Gardner et al.,



Figure 3.25: Comparison of approximate sampling methods. Mean and standard error on the mean of the log of the 2-Wasserstein distance estimation between true posterior distribution and approximate sampling distribution for three approximate sampling methods. We repeated the estimation for 32 functions sampled at random from a Gaussian process with zero mean and ARD kernel, and for different numbers of training points. The reduced-rank kernel approximation (RRGP) is more accurate compared to sparse-spectrum approximation (SSGP). Moreover, the decoupled-based approximation is more accurate than the weight-space approximation.

2018) and PyTorch. It benefits from the focus on scalability of GPytorch, and easy integration with the Pyro probabilistic programming language for advanced inference methods. It includes many advanced features such as multiobjective and multi-fidelity optimization, allowing to deal with cases where observations come from multiple information sources with different fidelity and/or cost. It can easily be used for the scenario with binary outputs.

Emukit is another Python package for Bayesian optimization, which includes multi-fidelity optimization (Paleyes et al., 2019) and provides support for preference learning and optimization.

The MATLAB code used in this work is available at https://github.com/TristanFauvel/ BO_toolbox. It contains, in particular, implementations for the approximate sampling methods.

3.8 Conclusion

In this work, we made the following contributions:

- we performed the first large-scale benchmarking of existing acquisition rules in both Binary Bayesian Optimization and Preferential Bayesian Optimization.
- we proposed an analytically tractable decomposition of uncertainty in GP classification into epistemic and aleatoric components. By interpreting exploration as epistemic uncertainty sampling, we proposed two new, highly efficient acquisition functions.
- we established a new state-of-the-art acquisition rule for Preferential Bayesian Optimization: the Maximally Uncertain Challenge acquisition (MUC).
- we established a new state-of-the-art acquisition rule for Binary Bayesian Optimization, which is a generalization of GP-UCB to models with probit likelihood (UCB_Φ).

- By using a simple decomposition of batch preference learning into independent pairwise comparisons, we generalized the Maximally Uncertain Challenge to batch preferential optimization, where it shows superior performance compared to state-of-the-art methods.
- By using recent developments in Gaussian process finite-rank approximation and sampling, we proposed efficient sampling techniques for sampling-based acquisition functions, which are especially relevant in batch optimization, when multiple comparisons are performed simultaneously (Sui and Burdick, 2017; Siivola et al., 2020).

Problem 3.1 is related to the best-arm identification in the bandits literature, and research in these fields cross-pollinated. Bayesian optimization algorithms are also referred to as Gaussian process bandits or kernelized bandits. However, in multi-armed bandits, the goal is to minimize the cumulative regret, whereas, in Bayesian optimization, the goal is to minimize the regret after spending the evaluation budget, which favors more exploratory behaviors (what is called best arm identification, or pure exploration, in the bandits literature (Bubeck et al., 2009)). Moreover, in Bayesian optimization, the input space is usually continuous. The bandit equivalent of PBO are dueling bandits (see Bengs et al. (2021) for a review), where at each iteration two arms are compared.

In general, theoretical convergence guarantees and regret bounds are still an open question in Bayesian Optimization. Srinivas et al. (2010) established the first sublinear regret bound for GP-UCB.Vakili et al. (2020) derived regret bounds in the noise-free setting with unknown kernel hyperparameters. However, to the extent of our knowledge, no such guarantees exist for non-Gaussian likelihoods.

Δ						
Random	0.09	0.09	0.15	0.03	0.12	0.00
BKG	0.06	0.18	0.18	0.06	0.00	0.18
TS	0.12	0.21	0.24	0.00	0.15	0.32
UCB_f	0.06	0.00	0.00	0.00	0.18	0.38
Binary EI	0.03	0.00	0.12	0.00	0.15	0.29
UCB_Φ	0.00	0.29	0.24	0.24	0.35	0.47
в						
Random	0.03	0.03	0.06	0.00	0.09	0.00
BKG	0.00	0.00	0.03	0.00	0.00	0.12
TS	0.09	0.00	0.03	0.00	0.18	0.32
UCB_f	0.12	0.06	0.00	0.18	0.35	0.50
Binary EI	0.09	0.00	0.00	0.06	0.32	0.47
UCB_Φ	0.00	0.06	0.00	0.26	0.38	0.59
	UCB_{Φ}	Binary EI	UCB_{f}	TS	BKG	Random

3.9 Supplementary

Figure 3.26: Detailed results of performance comparison between acquisition functions in Binary Bayesian optimization. Each entry (*i*,*j*) corresponds to the fraction of benchmarks functions for which *i* beats *j* according to the Mann-Whitney U test at $\alpha = 5 \times 10^{-4}$ significance based either on the best value found (**A**) or the Area Under the Curve (**B**). UCB_Φ largely outperforms other acquisition rules with respect to both criteria.

Δ									
Random	0.00	0.00	0.00	0.00	0.00	0.00	0.18	0.03	0.00
EI	0.00	0.00	0.06	0.03	0.03	0.06	0.18	0.00	0.18
EIIG	0.00	0.00	0.00	0.06	0.00	0.06	0.00	0.06	0.09
KSS	0.06	0.00	0.00	0.00	0.00	0.00	0.41	0.21	0.32
D. Thompson	0.00	0.00	0.00	0.03	0.00	0.15	0.47	0.35	0.44
DTS	0.06	0.00	0.09	0.00	0.03	0.06	0.50	0.29	0.47
BEI	0.03	0.00	0.00	0.09	0.03	0.15	0.47	0.35	0.41
UCB	0.00	0.00	0.06	0.09	0.00	0.24	0.50	0.41	0.56
MUC	0.00	0.00	0.06	0.15	0.06	0.21	0.50	0.41	0.59
В									
Random	0.00	0.00	0.03	0.00	0.00	0.00	0.06	0.06	0.00
EI	0.00	0.00	0.00	0.00	0.00	0.00	0.06	0.00	0.03
EIIG	0.00	0.00	0.03	0.09	0.00	0.12	0.00	0.18	0.12
KSS	0.00	0.00	0.06	0.00	0.00	0.00	0.09	0.06	0.03
D. Thompson	0.00	0.00	0.06	0.12	0.00	0.12	0.15	0.15	0.26
DTS	0.00	0.00	0.03	0.00	0.00	0.00	0.18	0.06	0.12
BEI	0.00	0.00	0.00	0.18	0.03	0.21	0.24	0.18	0.35
UCB	0.00	0.00	0.06	0.21	0.00	0.24	0.24	0.26	0.29
MUC	0.00	0.00	0.03	0.15	0.00	0.21	0.26	0.21	0.32
	MUC	UCB	BEI	DTS	D. Thompson	KSS	EIIG	EI	Random

Figure 3.27: Detailed results of performance comparison between acquisition functions in Preferential Bayesian optimization. Each entry (*i*,*j*) corresponds to the fraction of benchmarks functions for which *i* beats *j* according to the Mann-Whitney U test at $\alpha = 5 \times 10^{-4}$ significance based either on the best value found (A) or the Area Under the Curve (B).

Bayesian optimization in adaptive contexts

In chapter 2, we compared two strategies to optimize retinal prosthetic encoders: preferencebased and performance-based optimization. A limitation of the performance-based optimization approach is the fact that, as the optimization progresses, the task may become too simple so that it is impossible to improve the encoder further. Alternatively, if the task is so difficult that even with the best encoder performance is at chance level, responses are not informative. As a consequence, the task had to be carefully calibrated beforehand by tuning the size of the letter in the tumbling E task to get on average a 50 % chance of success with a randomly chosen encoder. Even then, the optimization did not succeed. At the time when I designed and performed these experiments, I could not find solutions to this problem in the literature. This led me to consider the problem of Bayesian optimization in adaptive contexts, where the context of the task used to measure the objective function (the size of the letter E in our experiment) is adapted to gather maximally informative responses. As we will see in this chapter, this problem is general and may arise in a variety of situations.

4.1 Introduction

4.1.1 Contextual Bayesian optimization with binary outputs

Consider an experiment in which the goal is to maximize the performance of a system by tuning some configuration parameters. The system can either pass or fail a test, and the result is stochastic. Assume that the system can work in different contexts or regimes, for which the difficulty of the test varies. If the regime is such that the system always fails (or passes) the test for any configuration, then we are wasting our evaluation budget. By contrast, adaptively selecting the context in which the test is performed could improve optimization efficiency. Consider for example the problem of tuning the optics of a patient with visual refractive error, such as myopia or astigmatism. Ophthalmologists measure the ability of the patient to recognize optotypes such as Snellen letters when changing both the vergence of the optics (the parameters to optimize) and the size of the Snellen letters (the context). If the letters are too big or too small, the responses will be uninformative about the optimal correction. A similar problem may arise in many situations. For example, when tuning the gait parameters of a robot walking in environments with various ground textures. The robot can either fall or succeed in walking, and a single set of parameters should allow it to walk across different conditions. We denote this problem, in the Bayesian optimization framework, Bayesian optimization in adaptive contexts.

The problem of choosing informative contexts does not arise in standard optimization, since

the objective function is directly measured (possibly with homogeneous noise), and to the extent of our knowledge, it has not been considered previously.

A similar problem may arise when trying to learn the preference of people. As an example of contextual preference comparisons, consider the problem of optimizing signal processing in hearing aid devices based on patients' preference judgments (Nielsen et al., 2014). The goal is to find the best setting for the device, across any kind of auditory stimulus a user may experience in daily life (the contexts). To do this, Nielsen et al. (2014) asked participants to compare pairs of device settings when listening to excerpts of music. New pairs of settings were adaptively selected using preferential Bayesian optimization so as to efficiently learn the preferred one. However, for a given auditory stimulus, there may not be a significant preference for one candidate setting over another, so that the response of the subject is random. For example, if the auditory stimulus that is presented is of small intensity, the two candidate settings may lead to the same perceptual experience, whereas for louder sounds one of them may lead to distortion. As a consequence, adaptively selecting the stimulus and the candidate settings could improve sample-efficiency.

Intuitively, a good strategy would be to select new candidate settings that are promising and to evaluate these settings in a context in which the output of a query is expected to be as informative as possible. The selection of the setting is determined by a tradeoff between exploration and exploitation, as in standard Bayesian optimization, while the selection of an informative context is an active learning problem. In this chapter, we show how methods from Bayesian optimization and Bayesian active learning can be combined to efficiently optimize the objective, by adaptively choosing the context and the new setting to evaluate.

Initially, for simplicity, we consider the case where the context doesn't affect the maximum of the function being optimized. For example, in the previous visual example, this would imply that the ideal lens correction was the same, independent of the stimulus being presented (e.g. the size of the letters). In general, however, this assumption may not be true. For instance, in the application of preferential Bayesian optimization to the optimization of hearing aids, the best settings of the device may vary depending on the task (e.g. listening to music, or having a conversation),

4.1.2 Problem formulation

Consider a function g(x) that corresponds to the implicit value of a setting x. In the refractive error correction example, this function would be related to the level of blur induced by given optics parameters $x \in \mathcal{X}$, where \mathcal{X} is a bounded set denoted as the search space. We are interested in solving the global optimization problem:

$$\boldsymbol{x}^{\star} = \operatorname*{arg\,max}_{\boldsymbol{x}\in\mathcal{X}} g\left(\boldsymbol{x}\right). \tag{4.1}$$

We cannot directly measure g(x). We are interested in scenarios in which the information brought by a function evaluation about the optimum of g depends on the context. This arises for example if function evaluation is subject to heteroscedastic noise, but for simplicity, we will restrict ourselves to the case where the result of an evaluation c is binary, which is known as binary BO (Tesch et al., 2013). Rather, we can only measure whether the system passes or fails a test in a given context *s* (whether the patient correctly identifies a letter of a given size for example). As is standard in Bayesian optimization with binary outcomes (Tesch et al., 2013), we assume that the outcome *c* of any evaluation is binary, such that:

$$c \sim \mathcal{B}\left(\pi\left(f\left(\boldsymbol{s},\boldsymbol{x}\right)\right)\right),\tag{4.2}$$

where \mathcal{B} represents the Bernoulli law, π is an inverse link function, typically the normal cumulative distribution function Φ .

To begin with, we will assume that there exists a configuration x that is the best across all contexts. That is, the maximum of $f(s, \cdot)$ does not depend on s. To come back to our example on tuning the optics of a patient with refractive error performing a visual acuity test, it is clear that the optics parameters x that minimize the blur will maximize performance (the probability of correctly identifying a letter) for any size of letter s. This condition can be formalized as:

$$\forall s \in \mathcal{S}, \underset{x \in \mathcal{X}}{\operatorname{arg\,max}} g\left(x\right) = \underset{x \in \mathcal{X}}{\operatorname{arg\,max}} f\left(s, x\right), \tag{4.3}$$

where S is a bounded set denoted as the context space.

The context determines, for a given x, the probability of success. In our refractive error correction experiment, choosing s can be seen as tuning the difficulty of the test by varying the letter size or contrast.

The function f directly determines the amount of information that a given context brings about g(x). To see this, consider as an example the case where we want to optimize a system that can either pass or fail a test, with success probability:

$$P(s, x) = \Phi(h(s)g(x) + b), \qquad (4.4)$$

where *b* is a constant. In this example, we assume that *g* and *h* are non-negative. $\Phi(b)$ thus corresponds to the baseline success rate (in our example on refractive error correction, this would be the probability of success if the patient answers randomly).

For any *x* such that g(x) > 0, for large values of h(s), the probability will saturate with $P(s, x) \approx 1$. In our refractive error correction example, this would correspond to a task that is too easy because the size of the letter is too large. By contrast, is h(s) = 0, $P(s, x) \approx \Phi(b)$ for any *x* (the task is too difficult and responses are random). Ideally, for any *x*, one would like to select a context where responses are informative about the underlying objective *g*. Our problem is thus to find a sampling rule to select *x* and *s* at each iteration so as to find the maximum of *g* as efficiently as possible.

As is standard in Bayesian optimization, in the following sections, we will build a surrogate model of *f* by putting a zero-mean Gaussian process prior on it, that is $f \sim GP(0, k(., .))$, where *k* is a kernel.

The reported solution to problem 3.1 after spending the evaluation budget is the maximum of the posterior mean for an arbitrary value of *s*. This corresponds to the standard assumption that the decision-maker is risk-neutral, i.e. decisions are based on the expected value, without considering uncertainty.

4.2 Acquisition rules for Bayesian optimization in adaptive contexts

4.2.1 Knowledge gradient acquisition rule

A standard acquisition rule when function evaluations are noisy is the knowledge gradient (KG), introduced by Frazier et al. (2009). The knowledge gradient is the simple idea of measuring the alternative that produces the greatest value, in expectation, from a single observation.

Given that the reported solution at time *t* is the maximum of the posterior mean of the Gaussian process μ_t^* after observing \mathcal{D}_t , if we can afford one more evaluation $(x_{t+1}, c_{t+1}), \mu_{t+1}^* - \mu_t^*$ corresponds to the change in expected value of the solution. The knowledge gradient corresponds to the expected increase in the posterior mean:

$$\operatorname{KG}(\boldsymbol{x}) = \mathbb{E}_{c \sim p(c|\mathscr{D}_t)}(\mu_{t+1}^{\star} - \mu_t^{\star}|\mathscr{D}_t, \boldsymbol{x}_{t+1} = \boldsymbol{x}), \tag{4.5}$$

where μ_t^* is the maximum of the posterior mean after observing $\mathscr{D}_t = (X, c_{1...,t})$, and μ_{t+1}^* is the maximum of the posterior mean after observing (\mathscr{D}_t , (\mathbf{x} , c)). The knowledge gradient acquisition rules selects the point maximizing KG at each iteration.

Previously, the knowledge gradient has been applied to standard BO with continuous outputs (Frazier et al., 2009), and to binary BO with linear Bayesian classifier as a surrogate model (Wang et al., 2016). Here we generalize the knowledge gradient to Gaussian process classification and preference learning. In particular, we show how to approximate its gradient for efficient optimization. With binary outputs and contexts, the knowledge gradient can be expressed as:

$$KG(s, x) = \mu_c(s, x)(\mu_{1,t+1}^{\star} - \mu_t^{\star}) + (1 - \mu_c(s, x))(\mu_{0,t+1}^{\star} - \mu_t^{\star}),$$
(4.6)

where $\mu_c(s, x) = P(c = 1 | x, \mathcal{D}_t) = \mathbb{E}_f[\Phi(f(s, x)) | \mathcal{D}]$ is the predictive class distribution, and $\mu_{1,t+1}^{\star}$ (resp. $\mu_{0,t+1}^{\star}$) is the maximum of the posterior mean for an arbitrary context s_0 after observing $(\mathcal{D}, (x, 1))$ (resp. $(\mathcal{D}, (x, 0))$). That is:

$$\mu_{c,t+1}^{\star} = \max_{\mathbf{x}' \in \mathcal{X}} \mathbb{E}_f[f(\mathbf{s}_0, \mathbf{x}') | \mathcal{D} \cup (\mathbf{s}_0, \mathbf{x}, c)].$$

$$(4.7)$$

Figure 4.1 shows an example of contextual binary knowledge gradient on a simple function.

The gradient of the knowledge-gradient, required for efficient maximization of the acquisition function, can be computed when using the Laplace approximation (see appendix E). However, this acquisition function can be impractically slow for problems with moderately high dimensions. Indeed, computing the knowledge gradient requires maximizing the latent mean function for the two possible outputs of the next query. To overcome this problem, we looked for another acquisition rule that would be faster, in particular for applications in human-in-theloop optimization.

4.2.2 Sequential acquisition

Given that the maximum of $f(s, \cdot)$ does not depend on the context s (condition 4.3), intuitively, the value of a new candidate x_{t+1} in terms of exploration-exploitation trade-off does not depend on s.



Figure 4.1: Contextual binary Knowledge Gradient. A. Example of binary task where the function f is of the form f(s, x) = h(s)g(x). In this example, the probability of correct response is 0.5 for any point (s, x) such that s = 1, so that responses are uninformative in this context. B. Predictive class distribution $\mu_c(s, x)$ for a Gaussian process classification model with a kernel of the form $k((s, x), (s', x')) = k_1(s, s')k_2(x, x')$. C. Corresponding contextual binary knowledge gradient.

As a consequence, a possible heuristic would be to select x_{t+1} to trade-off exploration and exploitation, as in standard Bayesian optimization, then to select s_{t+1} so as to gather responses as informative as possible about $f(\cdot, x_{t+1})$.

The first step of selecting x_{t+1} could be made using standard acquisition rules for binary Bayesian optimization (Tesch et al., 2013). Once the new candidate x_{t+1} is chosen, selecting a context in which the response is as informative as possible about $f(\cdot, x_{t+1})$ is an active learning problem (Houlsby et al., 2011).

In this work, we use Gaussian Process Upper Credible Bound (GP-UCB) (Srinivas et al., 2010) and Thompson Sampling (Thompson, 1933) to choose x_{t+1} . Thompson sampling consists in sampling x_{t+1} according to its posterior probability of being the maximum, that is we sample it from $p(x^*|\mathscr{D}_t)$. It shows competitive performance in standard Bayesian optimization (see e.g. Basu and Ghosh (2017)) as well as its generalizations to Preferential Bayesian Optimization (Gonzalez et al., 2017; Sui et al., 2017). The method we use to sample from $p(x^*|\mathscr{D}_t)$ is described in details in 3.6.1. Briefly, following Hernández-Lobato et al. (2014) we first draw a sample from a finite-dimensional approximation to the posterior GP, then maximize this sample to find x_{t+1} . To draw the sample from the posterior GP, we used the decoupled-bases approximation by Wilson et al. (2020a) jointly with the reduced-rank approximation described in Solin and Särkkä (2020).

We use the following generalization of GP-UCB to the case of binary ouputs (introduced in chapter 3):

$$\alpha_{\text{UCB}}(\boldsymbol{x}) = \mathbb{E}(\Phi(f(\boldsymbol{x})) + \beta \sqrt{\mathbb{V}(\Phi(f(\boldsymbol{x})))}$$
(4.8)

The term $\mathbb{V}(\Phi(f(\mathbf{x})))$ and its gradient can be computed analytically (see appendix D). We chose $\beta = \Phi^{-1}(0.95)$.

Once, x_{t+1} is chosen, the idea is to select a context s_{t+1} such that the query will be as informative as possible about the underlying function $f(\cdot, x_{t+1})$. This can be seen as a Bayesian active learning problem. To select the context conditionally on x_{t+1} , we used Bayesian Active Learning by Disagreement (BALD, Houlsby et al. (2011)). It consists in selecting points that maximize mutual information between the response and the underlying latent function. That

is, for a given x_{t+1} , the context s_{n+1} is chosen by maximizing:

$$I(c, f | \mathscr{D}, \mathbf{s}_{n+1}, \mathbf{x}_{t+1}) = H(y | \mathscr{D}, \mathbf{x}_{t+1}, \mathbf{s}_{n+1}) - \mathbb{E}_f [H(y | \mathscr{D}, \mathbf{x}_{t+1}, \mathbf{s}_{n+1}, f)],$$
(4.9)

with respect to s_{n+1} . This acquisition criterion can be efficiently approximated and maximized. We termed these sequential acquisition rules Thompson sampling with active learning by disagreement (TS-ALD) and GP-UCB with active learning by disagreement (UCB-ALD). Although these heuristics are less principled than the Knowledge Gradient which jointly select *s* and *x*, they are much faster since they only require to solve one optimization problem.

4.2.3 Generalization to preference judgements

We next considered how to extend our method to preferential optimization, where outputs are comparisons between pairs of inputs (x, x') (also called 'duels'), of the form $x \succ x'$ or $x' \succ x$.

In preferential Bayesian optimization, a preference function *h* determines the outcome of a duel between *x* and *x'*, with $P(x \succ x') = \mathcal{B}(\pi(h(x, x')))$ (Brochu et al., 2010b; Gonzalez et al., 2017). *h* corresponds to the difference between the latent values of the points, described by function *g*. The goal is to find the maximum of the latent value *g*.

In common with the previous section, we assume that the value function *g* cannot be accessed directly, but rather via a function *f*, that depends on the context, so that the response in a duel (*x*, *x'*) in context *s* is given by $c \sim \mathcal{B}(\pi (f(s, x) - f(s, x')))$.

Again, we assume that the contextual value function *f* is such that:

$$\forall s \in \mathcal{S}, \underset{x \in \mathcal{X}}{\operatorname{arg\,max}} g\left(x\right) = \underset{x \in \mathcal{X}}{\operatorname{arg\,max}} f\left(s, x\right). \tag{4.10}$$

Using this formalism, the generalization of the knowledge gradient to this scenario is straightforward. Sequential acquisition rules can also be generalized by selecting the duel (x, x') using standard methods, and then choosing the context of the duel *s* using BALD. The generalization of Thompson sampling to the preference scenario is Kernel Self-Sparring (KSS) (Sui and Burdick, 2017). In KernelSelfSparring, the duels are made of two samples from $p(x^*|\mathcal{D}_t)$. We termed the generalization of TS-ALD to the preference scenario KSS-ALD. Similarly, we generalize UCB-ALD to the case of preference judgments by using the Maximally Uncertain Challenge acquisition rule instead of UCB (see chapter 3 for details on MUC).

4.3 Experiments with synthetic test functions

To evaluate the performance of our method, we ran optimization experiments using a set of 34 functions (Surjanovic and Bingham, 2021). We inferred the corresponding hyperparameters for three different kernels (squared exponential, Matérn 3/2, and Matérn 5/2) using maximum likelihood estimation with 1000 randomly chosen samples. We then determined for each function the kernel that was better at describing the function by measuring the root-mean-squared error on 3000 points and used this kernel in subsequent optimizations. The benchmark functions are listed in appendix J. During the optimization experiments, we performed Gaussian process classification using the Laplace approximation (see I).

We introduced a scalar context variable $s \in [0, 1]$, so that for a test function g the response of a system query is 1 with probability $P(s, x) = \Phi(sg(x))$ and 0 otherwise.

To take the context variable into account when building the GP surrogate model, we used the following kernel: for a base kernel *k* determined using the aforementioned procedure, we used: $k_s((\mathbf{x}, s), (\mathbf{x'}, s')) = ss'k(\mathbf{x}, \mathbf{x'})$

To compare the different algorithms, we used the stratified analysis method proposed by Dewancker et al. (2016). Briefly, for each benchmark function, we performed pairwise comparisons between acquisition functions using the Mann-Whitney U test at $\alpha = 5 \times 10^{-4}$ significance on the best value found at the end of the optimization sequence. This determines a partial ranking based on the number of wins.

Ties are then broken by running the same procedure, but based on the Area Under Curve, which is related to the speed at which the algorithm reaches the optimum. This generates a new partial ranking, based on which a Borda score (Dwork et al., 2001) is attributed to each acquisition function (the Borda score of a candidate is the number of candidates with a lower rank). Then, rankings from different benchmarks are aggregated by summing the Borda scores to establish a global ranking. This can be seen as a weighted vote from each benchmark function. Results are presented on table 4.1.

4.3.1 Binary Bayesian optimization in adaptive contexts

We then ran optimizations on each function for 40 different random number generator seeds, on 60 iterations. We compared the different acquisition rules: the contextual binary knowledge gradient (cBKG), UCB-ALD and TS-ALD, with the following controls:

- Fully random, in which (*s*, *x*) was chosen at random at each iteration,
- *x* chosen using TS and *s* chosen randomly,
- *x* chosen using UCB and *s* chosen randomly,
- (*s*, *x*) selected using BALD.

To avoid saturation effects when transforming the benchmark functions through the nonlinearity, we scaled the functions so that they have mean 0 and variance 1. Examples of regret curves are presented in figure 4.2, and the results of the stratified analysis are summarized in table 4.1. UCB-ALD and TS-ALD show superior performance compared to cBKG (4.1), despite the fact that, compared to cBKG, where x and s are jointly selected, these sequential decision strategies induce an adaptivity gap (Jiang et al., 2019). Here, we used the Laplace approximation, given that Expectation Propagation (EP) is known to improve performance compared to the Laplace approximation and that the gradient of cBKG with EP seems intractable, this suggests that TS-ALD and UCB-ALD are better acquisition functions in general. The three heuristics outperformed the different controls.

4.3.2 Preferential Bayesian optimization in adaptive contexts

We then evaluated the performance of our sequential acquisition rules in their generalization to the case of preferential judgments: KSS-ALD and MUC-ALD. Given the limited performance of BKG in the previous experiments and the fact that the preferential version of Knowledge Gradient is impractically slow, we did not evaluate this acquisition function.



Figure 4.2: Results across 5 different objective functions and 5 different methods in binary Bayesian optimization with adaptive contexts. Solid lines correspond to the mean value of the inferred maximum across the 40 repetitions. Shaded areas correspond to standard error on the mean. Either cBKG, UCB-ALD, or TS-ALD outperform the other acquisition rules, however, there is a lot of variability from one experiment to another. The cBKG acquisition rule tends to be less robust than the sequential heuristics.

Acquisition rule	Rank	Borda score
Upper Credible Bound with Active Learning by Disagreement	1	144
Thompson Sampling with Active Learning by Disagreement	2	83
Binary Knowledge Gradient	3	59
Upper Credible Bound (random context)	4	46
Bayesian Active Learning by Disagreement	5	35
Thompson Sampling (random context)	6	19
Random	7	10

Table 4.1: Comparison of acquisition functions on benchmarks for cBBO. UCB-ALD largely outperforms the other rules for which the context (TS, UCB, random) or the correction (BALD, random) are not adaptively selected. Importantly, the sequential acquisition rules (TS-ALD and UCB-ALD) outperform cBKG.

We then ran optimizations on each test function for 40 different random number generator seeds, on 60 iterations. We compared KSS-ALD and MUC-ALD to the following controls:

- Fully random, in which (s, x_1, x_2) was chosen at random at each iteration,
- *x* chosen using KSS and *s* chosen randomly,



Chapter 4. Bayesian optimization in adaptive contexts

Figure 4.3: Results across 5 different objective functions and 5 different methods in preferential Bayesian optimization with adaptive contexts. Solid lines correspond to the mean value of the inferred maximum across the 40 repetitions. Shaded areas correspond to standard error on the mean. Sequential acquisition rules (KSS-ALD and MUC-ALD) consistently outperform the others.

- *x* chosen using MUC and *s* chosen randomly,
- (s, x_1, x_2) selected using BALD.

Examples of regret curves are shown in figure 4.3, and the results of the stratified analysis are summarized in figure 4.2. Again, adaptive selection of the context leads to superior performance compared to controls.

Acquisition rule	Rank	Borda score
MUC-ALD	1	99
KSS-ALD	2	37
Kernel Self-Sparring	3	22
Bayesian Active Learning by Disagreement	4	9
Random	5	5

Table 4.2: Comparison of acquisition functions on benchmarks for cPBO. The sequential acquisition rules outperform the controls where the context (KSS, random) or the correction (BALD, random) are not adaptively selected. MUC-ALD is the most efficient acquisition function.

4.4 Application: adaptive optimization in psychometric measurements

We applied our method to the problem of optimizing the lens parameters in a subjective refractive error measurement task. This is a good example to illustrate the application of the new methods we propose as it is intuitive and grounded in a well understood perceptual phenomenon. Moreover, developments of our approach to this specific application may prove useful in clinical applications in the near future.

In patients with refractive error, the pathological shape of the eye causes the visual image to focus outside of the retina surface, inducing blur. With myopia, the eye is too large compared to its focusing power, which induces spherical refractive error. In astigmatism, the cornea is ellipsoid instead of spherical, causing a cylindrical refractive error.

Half of the world's population is expected to be myopic (short-sighted) by 2050 (Holden et al., 2016). When trying to find the optimal correction for a patient with refractive error, the ophthalmologist changes both the vergence of the optics (the parameter to optimize) and the size of the optotypes (standardized figures used to measure visual acuity, such as the Snellen letters) in standard recognition tasks. If the letters are too big or too small, the responses will be uninformative about the optimal correction. The usual correction step is 0.25 diopters, but recent technologies, such as Essilors' Advanced Vision Accuracy (Essilor International, Charenton-le-Pont, France) allow for 0.01 diopters precision. Moreover, many patients combine several types of refractive errors, such as myopia and astigmatism, which requires tuning several parameters of the lens at the same time. As a consequence, subjective refractive error measurements need to be adapted to be rapid and precise.

We evaluated Bayesian optimization in adaptive contexts in a simulation of a patient's subjective refractive error measurement. To do so, we first built a model of the response in a visual task where the patient has to identify an optotype displayed on a screen.

4.4.1 Model of patients' responses

Consider a subject performing an n-alternatives forced-choice visual task, such as identifying a letter on a screen. The size of the letter, s (measured in log of the visual angle, in minutes of arc), determines the difficulty of the task. For a given optical correction x, the probability of correct response is well described by a sigmoid function, called the psychometric function (see e.g. Fülep et al. (2019) or Tokutake et al. (2011)):

$$P(c = 1|s) = \gamma + (1 - \gamma)\Phi(as + b),$$
(4.11)

where γ is the correct response chance level, *a* and *b* are respectively denoted as the slope and the intercept. *a* is assumed to be strictly positive.

For a patient with optimal correction, the visual acuity is usually defined as the value of *s* at which $P(c = 1|s) = \frac{1+\gamma}{2}$, or equivalently the inflexion point of the curve. So, the visual acuity is related to the parameters of the psychometric curve by $VA = -\frac{b}{a}$ (see figure 4.5).

For lens parameters, *x*, the response of the patient is thus determined by:

$$P(c = 1|s, x) = \gamma + (1 - \gamma)\Phi(a(x)s + b(x)).$$
(4.12)

We introduce a function *f* such that: $P(c = 1 | s, x) = \Phi(f(s, x))$. By definition:

$$f(s, \mathbf{x}) = \Phi^{-1}[\gamma + (1 - \gamma)\Phi(a(\mathbf{x})s + b(\mathbf{x}))]$$
(4.13)

Experimental evidence show that the psychometric curve flattens (the slope a(x) decreases) and the intercept b(x) decreases as the blur induced by defocus increases (Carkeet et al., 2001). As a consequence, a(x) and b(x) share the same optimum x^* which corresponds to minimal blur. Moreover, visual acuity get worse as blur increases, as a consequence, x^* is the minimum of $VA(x) = -\frac{b(x)}{a(x)}$. The problem of improving visual acuity by minimizing $VA(x) = -\frac{b(x)}{a(x)}$ thus satisfies condition 4.3, as for any *s*, the maximum of *f* is the minimum of *VA*.

A simple formula relates the spherico-cylindrical error to image blur, and hence visual acuity (Raasch, 1995; Blendowske, 2015): by noting *S* the spherical refractive error and *C* the cylindrical refractive error, measured in diopters, the blur is given by:

$$\beta = \sqrt{\frac{1}{2}(S^2 + (S+C)^2)} \tag{4.14}$$

Visual acuity can then be related to blur using the formula: $MAR_{rel} = 1 + \beta^2$ (Blendowske, 2015), where MAR_{rel} is the ratio between the Minimum Angle of Resolution (MAR, measured in minutes of arc) and the MAR for the best correction. A visual acuity of 20/20 corresponds to logMAR = 0, so we considered that the MAR for the best correction is 0.

In our model, we assumed that the slope was constant, and only the intercept varied with correction parameters x. We repeated the experiment for slopes from 1.0 logMAR⁻¹ to 8.0 logMAR⁻¹.

4.4.2 Gaussian process model

The use of Gaussian process classification for psychometric function estimation has been introduced by Gardner et al. (2015); Song et al. (2017, 2018) in the context of audiometry. We follow the same approach and build a Gaussian process model of the subject's responses in the task.

For γ large enough, we have $P(c = 1|s) \approx \Phi(a(x)s + b(x))$. Since in practice, with a letter chart, $\gamma^{-1} = n = 26$, we will make this simplifying assumption in building our surrogate model.

We put a Gaussian process prior on f, with zero mean and kernel k_{ψ} defined as:

$$k_{\psi}((s, x), (s', x')) = \theta ss' + k(x, x')$$
(4.15)

This kernel reflects the structure of the function *f* (see equation 4.13) with the assumption $\gamma \approx 0$ (see supplementary 4.6.2.

We used a squared exponential kernel as k. To minimize the effect of kernel hyperparameters, we simulated the response to 1000 random pairs (s, x) and inferred the kernel hyperparameters using maximum likelihood estimation.

4.4.3 Results

We repeated the simulated experiment 20 times for 260 iterations, for 8 different slopes values (evenly spaced between 1.0 logMAR⁻¹ and 8.0 logMAR⁻¹). The search space was $\mathcal{X} =$



Figure 4.4: A. Visual acuity as a function of the correction parameters S (spherical correction) and C (cylindrical correction) in diopters (δ). The slope of the psychometric curve is kept fixed at 5.0 logMAR⁻¹. B. Probability of correct response in a 26-alternatives forced choice task for a stimulus of visual angle 1'. In blue regions, the subject will always perform at chance level, independently of the correction parameters, so responses will be uninformative.

 $[-4\delta, 4\delta] \times [-4\delta, 4\delta]$ and the context space $S = [-1\delta, 2\delta]$. Regret curves in the experiments were the slope of the psychometric function was set at 5.0 logMAR⁻¹ are shown on figure 4.5, and the results of the stratified analysis are summarized in table 4.3.

Both UCB-ALD and TS-ALD led to rapid and consistent improvement of the visual acuity closed to its optimal value of logMAR = 0. UCB-ALD is the best performing algorith (see table 4.3), with a mean visual acuity at the end of the optimization sequence of 2.48×10^{-2} logMAR in the case of a slope of $5.0 \log MAR^{-1}$ (s.e.m = 8.22), 4.5). At the end of the optimization sequence, the an average spherical correction error of $2.74 \times 10^{-2}\delta$ (s.e.m. = 5.36×10^{-2}). The average cylindrical correction error is $1.53 \times 10^{-1}\delta$ (s.e.m. = 8.22×10^{-2}).

However, when using random sampling, the optimization consistently failed (figure 4.5). Overall, the results are consistent with the one obtained on the synthetic benchmarks (table 4.3).

Here, we considered that the objective is a black box. However, for a specific application in a clinical setting, parameterizing the objective using domain knowledge would likely lead to faster and more robust improvement (see e.g. (Cox and De Vries, 2017)).
Chapter 4. Bayesian optimization in adaptive contexts



Figure 4.5: *A.* Psychometric curves for various refractive errors corresponding to different visual acuities, with a slope of the psychometric function set at $5.0 \log MAR^{-1}$. The baseline corresponds to the chance level in the n-alternatives forced-choice task (here n = 26). Visual acuity of 20/20 corresponds to $\log MAR = 0$. *B.* Visual acuity with the inferred best parameters throughout the optimization. The lines correspond to the average over 60 repetitions, whereas the shaded areas correspond to standard error on the mean. UCB-ALD and TS-ALD rules both lead to faster convergence towards optimal correction ($\log MAR = 0$) compared to controls where the context (TS, UCB, random) or the correction (BALD, random) are not adaptively selected.

Acquisition rule	Rank	Borda score
Upper Credible Bound with Active Learning by Disagreement	1	35
Thompson Sampling with Active Learning by Disagreement	2	26
UCB (random context)	3	15
Bayesian Active Learning by Disagreement	4	8
Thompson Sampling (random context)	5	3
Random	6	0

Table 4.3: Comparison of acquisition functions on the refractive error correction experiment with 8 different slopes (evenly spaced between between 1.0 $\log MAR^{-1}$ and 8.0 $\log MAR^{-1}$). The two heuristics combining Bayesian optimization and active learning (UCB-ALD and TS-ALD) outperform the controls where the context (TS, UCB, random) or the correction (BALD, random) are not adaptively selected.

4.5 Discussion

4.5.1 Summary of contributions

In this chapter, we introduced a new framework for Bayesian optimization: Bayesian optimization in adaptive contexts, in which the experimenter can choose the context where each function evaluation is performed. We proposed an acquisition function to jointly select the optimization variable and the context and showed how to generalize existing acquisition functions to this framework by combining them with Bayesian active learning. We show on synthetic benchmarks that the proposed algorithms outperform controls in which the context is not adaptively selected, in both Bayesian optimizations with binary or preference feedbacks. Finally, we showcase our framework on a concrete problem: tuning the parameters of the lens for patients with refractive errors.

4.5.2 Generalization to ConBO and Max-Min BO

So far, we considered the case where the optimum of the function f is the same across all contexts. In general, the maximum of $f(s, \cdot)$ may depend on s. In that case, several criteria may be considered. Either we may want to maximize the averaged performance of the system across contexts, or we may want to maximize the worst-case performance of the device. The former scenario corresponds to the ConBO framework (Pearce et al., 2020), whereas the second corresponds to Max-Min Bayesian optimization (Weichert and Kister, 2020). Our approach is a generalization of these frameworks as we consider cases in which the context directly determines the amount of information about f at the sampled points.

The generalization of the ConBO framework to the binary case is given by:

$$g(\mathbf{x}) = \mathbb{E}_{\mathbf{s}}\left[\pi[f(\mathbf{x}, \mathbf{s})]\right],\tag{4.16}$$

whereas the generalization to max-min optimization is:

$$g(\boldsymbol{x}) = \min_{\boldsymbol{s} \in S} \pi[f(\boldsymbol{s}, \boldsymbol{x})]. \tag{4.17}$$

We let the study of these generalizations for future work.

4.5.3 Related problems

Our work is related to Bayesian optimization with environmental variables (Klein et al., 2017; McLeod et al., 2017): in this scenario, some environments correspond to cheap but unreliable evaluations, and others to accurate but expensive evaluations. Our framework differs as we do not consider the cost of the queries, and we do not know a priori how informative a given context is.

Similarly, in the bandit setting, contextual GP bandit (Krause and Ong, 2011) extends the Gaussian process bandit framework to scenarios in which at each round corresponds a given random context, and actions are taken knowing this context. The goal is to identify the best arm averaged over all possible contexts. This framework has been generalized to dueling bandits (Dudík et al., 2015). Our work differs as we consider the problem of both selecting a context and action at each iteration.

The methods most related to our work are Bayesian optimization with expensive integrands (Toscano-Palmerin and Frazier, 2018), conditional BO (Pearce et al., 2020), and contextual policy search (Pinsler et al., 2019), where the goal is to maximize the expected performance conditionally on a random context variable. During the optimization process, this context variable is controlled by the experimenter. In Bayesian optimization with expensive integrands (Toscano-Palmerin and Frazier, 2018), the goal is to maximize a function averaged over a set of random conditions, when function evaluation is noise-free or with i.i.d. Gaussian noise. The authors used an acquisition rule that jointly selects the optimization variable and the context, which

is a generalization of the knowledge gradient. Pearce et al. (2020) further generalized by considering the case where function evaluations are performed in parallel and batch of inputs are selected at each iteration.

Our work is a generalization to the binary and preferential feedback scenarios, which implies new challenges. Indeed, even when the objective function is the same across different contexts, the amount of information we get from each evaluation varies with the context in a way that is a priori unknown. This problem arises in general when the amount of noise in function evaluation varies with the context, but we only considered the case of probit likelihoods for simplicity. We also introduced sequential acquisition rules that allow the generalization of acquisition rules used in Bayesian optimization.

Thatte et al. (2017) applied the Predictive Entropy Search acquisition rule (Hernández-Lobato et al., 2014) to PBO. The Predictive Entropy Search with Preferences (PES-P) samples comparisons that optimally reduce the uncertainty in the distribution of the objective function maximum with the least number of queries and it could in theory be applied in the contextual scenario. However, this acquisition function requires several approximations and is slow to compute.

We did not consider whether sequential acquisition also improves performance in the case of Gaussian likelihood, and this should be the topic of future work. Potential future developments also include the generalization of our method to batches, where several points are selected and evaluated simultaneously.

4.5.4 Code availability

Matlab implementation available at https://github.com/TristanFauvel/Contextual_BBO.

4.6 Chapter supplementary

Α .							
Random	0.03	0.00	0.00	0.03	0.06	0.00	0.00
TS	0.00	0.00	0.03	0.03	0.06	0.00	0.06
BALD	0.06	0.00	0.09	0.06	0.00	0.03	0.12
UCB	0.00	0.00	0.06	0.00	0.12	0.06	0.24
cBKG	0.03	0.00	0.00	0.15	0.15	0.15	0.18
TS-ALD	0.03	0.00	0.09	0.09	0.29	0.12	0.32
UCB-ALD	0.00	0.12	0.24	0.03	0.32	0.44	0.50
B Random	0.00	0.00	0.03	0.00	0.09	0.00	0.00
TS	0.00	0.00	0.06	0.00	0.06	0.00	0.06
BALD	0.00	0.00	0.12	0.06	0.00	0.09	0.15
UCB	0.00	0.00	0.09	0.00	0.06	0.09	0.15
cBKG	0.03	0.00	0.00	0.06	0.21	0.12	0.09
TS-ALD	0.00	0.00	0.12	0.03	0.26	0.09	0.29
UCB-ALD	0.00	0.24	0.38	0.09	0.65	0.47	0.71
	UCB-ALD	TS-ALD	cBKG	UCB	BALD	TS	Random

4.6.1 Supplementary results

Figure 4.6: Detailed results of performance comparison between acquisition functions in contextual Binary Bayesian optimization. Each entry (*i*,*j*) corresponds to the fraction of benchmarks functions for which *i* beats *j* according to the Mann-Whitney U test at $\alpha = 5 \times 10^{-4}$ significance based either on the best value found (A) or the Area Under the Curve (B).

4.6.2 Kernel derivation for the psychophysics experiment

For lens parameters, *x*, the response of the patient is determined by:

$$P(c = 1|s, x) = \gamma + (1 - \gamma)\Phi(a(x)s + b(x)).$$
(4.18)

For γ large enough, we have $P(c = 1|s) \approx \Phi(a(x)s + b(x))$. Since in practice, with a letter chart, $\gamma^{-1} = n = 26$, we will make this simplifying assumption in building our surrogate model.

We put a Gaussian process prior on f, with zero mean and kernel k_{ψ} defined as:

$$k_{\psi}((s, \boldsymbol{x}), (s', \boldsymbol{x}')) = \theta ss' + k(\boldsymbol{x}, \boldsymbol{x}')$$

$$(4.19)$$

This kernel reflects the structure of the function *f* (see equation 4.13) with the assumption $\gamma \approx 0$. Indeed, we have:

$$Cov(f(s, x), f(s', x')) = Cov(a(x)s + b(x), a(x')s' + b(x'))$$

= ss'Cov(a(x)s + b(x), a(x')s' + b(x'))
+ (s + s')Cov(a(x), b(x')) + Cov(b(x), b(x')) (4.20)

Α					
Random	0.06	0.00	0.00	0.03	0.00
BALD	0.09	0.00	0.00	0.00	0.06
KSS	0.09	0.03	0.00	0.15	0.24
KSS-ALD	0.09	0.00	0.06	0.21	0.26
MUC-ALD	0.00	0.47	0.44	0.53	0.59
в					
Random	0.00	0.00	0.00	0.03	0.00
BALD	0.00	0.00	0.00	0.00	0.03
KSS	0.00	0.00	0.00	0.12	0.12
KSS-ALD	0.06	0.00	0.06	0.24	0.24
MUC-ALD	0.00	0.50	0.53	0.56	0.74
	MUC-ALD	KSS-ALD	KSS	BALD	Random

Chapter 4. Bayesian optimization in adaptive contexts

Figure 4.7: Detailed results of performance comparison between acquisition functions in contextual Preferential Bayesian optimization. Each entry (*i*,*j*) corresponds to the fraction of benchmarks functions for which *i* beats *j* according to the Mann-Whitney U test at $\alpha = 5 \times 10^{-4}$ significance based either on the best value found (A) or the Area Under the Curve (B).

We put Gaussian process priors on functions *a* and *b*, with kernels k_1 and k_2 respectively. since a(x) and b(x) are a priori independent conditionally on *x*, the second term on the right-hand side vanishes, so that:

$$Cov(f(s, x), f(s', x')) = ss'k_1(x, x') + k_2(x, x')$$
(4.21)

Since we assumed the slope to be constant at value $\sqrt{\theta}$, this further simplifies to:

$$\operatorname{Cov}(f(s, \boldsymbol{x}), f(s', \boldsymbol{x}')) = \theta ss' + k(\boldsymbol{x}, \boldsymbol{x}')$$
(4.22)

5

5.1 Summary of contributions

5.1.1 Validation of a human-in-the-loop optimization strategy for retina implant encoders in simulated prosthetic vision

Model-based encoder

In this thesis, we provided a proof of principle for a retina implant encoder optimization strategy. Our approach was based on two main ideas. The first one is the parameterization of the encoder using approximate inversion of a patient-specific perceptual model mapping electrodes activity to percepts. We demonstrated on a previously published model predicting percepts experienced by patients equipped with retinal prosthetic devices, *pulse2percept* (Beyeler et al., 2017a, 2019b), that model-based parameterization of retina implant encoders leads to significant improvement in perceptual experience compared to encoders based on naive image downsampling in simulated prosthetic vision. However, the parameters of the encoder need to be fine-tuned in a patient-specific way. Indeed, previous work emphasized the variability of percepts experienced by patients equipped with prosthetic devices (Beyeler et al., 2019b). The second idea in our strategy is thus to use preferential Bayesian optimization to optimize the encoder parameters online, based on patient feedback via pairwise comparisons between encoders.

Encoder tuning with preferential Bayesian optimization

We choose Bayesian optimization since it is one of the most efficient approaches in optimizing black-box functions, meaning that this algorithm tends to rapidly find good solutions to global optimization problems (Turner et al., 2021). Among the many types of feedbacks we could gather from participants, we choose to query pairwise comparisons between encoders. This choice is motivated by the fact that objectively measuring the quality of a given encoder is non-trivial: a good encoder should allow perceiving high-level visual features for a wide range of stimuli and contexts. An alternative approach would have been to use an objective measure of visual performance. However, by contrast, this could lead to optimizing the encoder in a way that is not correlated with general visual function. As an extreme example, consider an encoder that converts a stimulus into a symbolic representation using visual Braille. In a reading task, such an integrated optical Braille recognition system would likely lead to good performance, without actually providing vision. Instead, we assumed that subjects implicitly attributed a

value to each encoder, taking into account many aspects of perceptions that a single task could hardly account for, and that this value could be inferred by pairwise comparisons.

We assessed our idea in healthy subjects, again making use of prosthetic vision simulation, and demonstrated that preference-based optimization is efficient in improving the prosthetic perceptual experience. Most importantly, this improvement was able to transfer to tasks related to visual acuity. These results in prosthetic vision simulation suggest that, by using models of prosthetic vision synergistically with preferential Bayesian optimization, it should be possible to significantly improve functional vision in patients with retina implants.

Robustness to model misspecification

A concern one may raise is that, in our experiments, the model on which the encoder was parameterized was the same as the model used to simulate vision, and only the parameters of the model were unknown. In real-life patients, the model used to parameterize the encoder may differ quite significantly from the true mapping from electrical pulses to percepts. This problem is recurrent, although not explicitly stated, in the literature on encoder optimization, where it is typically assumed that the model used to encode visual stimuli is in agreement with the model of the decoder (mapping stimulation to percepts). See e.g. Eckmiller et al. (1999); Nirenberg and Pandarinath (2012); Shah et al. (2019); Romeni et al. (2021) where this problem is apparent. As a consequence, the predicted performance of the optimized encoders can sometimes seem largely unrealistic (see Nirenberg and Pandarinath (2012) for a good example).

To address this caveat, we misspecified the model used to parameterize the encoder. By doing so, we ensured that even if the remaining free parameters were set at their true value, the perceptual quality was far worse than what could be achieved with an optimal encoder. However, we observed that the preference-based optimization procedure was still able to improve the quality of the encoder compared to randomly chosen parameters, meaning that the optimization was robust to a severe mismatch between the two models. However, the optimized encoders were worse than the optimal encoder, meaning that the optimization cannot compensate for the wrong parameterization. This can be interpreted as the fact that the misspecification restricted possible solutions to a subspace where the optimum was worse than the optimum in the whole space. This suggests that there is an important tradeoff between the flexibility of the encoder and the difficulty to optimize it.

Transfer learning for efficient and robust optimization

The surrogate model is the main component of a Bayesian optimization algorithm, and its choice (in particular, when using Gaussian processes, the choice of the kernel and its hyperparameters) can have dramatic effects on performance (see chapter 3, figure 3.2). For example, if the lengthscale of the kernel is too small, the model will not generalize as much as it could, which can lead to under-exploration. A difficulty we faced was the selection of an appropriate kernel and its hyperparameters when modeling the latent value function. This problem has received only little attention in the Bayesian optimization literature (Berkenkamp et al., 2019). The hyperparameters and the kernel are usually adapted over the course of the optimization, using Bayesian model selection (Rasmussen and Williams, 2006).

In pilot experiments, we followed this first approach and inferred the hyperparameters online, using type-II maximum-likelihood estimation. However, this has several drawbacks. First, this slows the acquisition process down, which conflicts with the criterion we had for our optimization strategy that the time between two queries would be much shorter than the query itself. Second, with a small amount of data, maximum likelihood estimation usually leads to overfitting, severely affecting the optimization success. To overcome these difficulties, we finally adopted the hyperparameter transfer introduced in chapter 2 and illustrated in figure 2.2. This strategy assumes that the correlation structure between the value of different encoders is roughly similar between different subjects. It showed good performance in practice, is easy to implement, and could be generalized to any optimization problem where a dataset related to the optimization problem exists. A more principled approach that could be used in the future would be to rely on a hyperprior on these hyperparameters. This hyperprior could be learned from previous experiments with other patients.

5.1.2 Uncertainty decomposition for efficient Bayesian optimization

In initial pilot experiments, we found that commonly used acquisition rules for preferential Bayesian optimization (Expected Improvement, Bivariate Expected Improvement, and Kernel-SelfSparring) did not perform well compared to random sampling. To understand why we performed synthetic experiments which confirmed the general limitations of these heuristics, as well as others (see chapter 3). By making use of recent developments in the field of Bayesian active learning, we developed a new acquisition strategy for Bayesian optimization with binary data that is fast and outperforms all existing acquisition rules in the binary, preferential, and batch preferential settings. This strategy is based on Bayesian uncertainty decomposition into aleatoric and epistemic uncertainty. Extracting the epistemic component of uncertainty allows to refine exploration and improve performance. We also improved several aspects of existing preferential Bayesian optimization algorithms, by introducing a specific preference kernel improving uncertainty estimation, and by leveraging recent developments in finite-rank kernel approximation and approximate sampling from Gaussian processes. We were then able to use these methods in our psychophysics experiments, where we observed a significant increase in performance.

5.1.3 Contextual optimization and limitations of performance-based optimization

One may argue that when using preference-based optimization, subjects may tend to favor aspects of the stimuli that could impede visual performance in daily tasks. For example, a subject may prefer a given encoder for subjective aesthetic reasons, unrelated to visual function. As a consequence, it could make more sense to directly optimize the encoder to maximize performance in visual tasks.

To investigate this idea, we compared the efficiency of preference-based optimization and performance-based optimization in tuning retina implant encoders. Instead of maximizing subject preference for encoders, we maximized an objective performance measure in a tumbling E task. By doing so, we highlighted a serious limitation of performance-based optimization with binary outputs, which is that the difficulty of the task has to be carefully tuned beforehand,

then adjusted throughout the optimization process.

This led us to consider the general problem of optimizing a system working in various regimes or contexts (the size of the letters in the tumbling E task) when the mutual information between the outcome of a function evaluation and the function value depends on the context. We only considered the case where feedbacks are binary, but this could also apply to a situation where function evaluation is subject to heteroscedastic noise. We approached this problem by combining acquisition strategies from Bayesian optimization (to select promising settings of the system) and Bayesian active learning (to select contexts in which function evaluations are expected to be informative). We showed, in particular, how this could be applied in visual psychophysics (chapter 4), and in the future, it could be applied to the problem of tuning retina implants.

Another issue for performance-based optimization, besides correctly setting the task difficulty, is ensuring that improvements in one task are able to transfer. Indeed, the optimization process could emphasize artifactual cues that increase the performance of the subject in the task, without overall improvement in percepts' quality. As a consequence, performance improvement may not correlate with better functional vision. This problem is well known in psychophysics, where overtraining can lead to absurd results. Ng and Liem (2018), for example, showed that dichromat subjects were able to significantly improve their performance in a color vision test by using spurious brightness cues. Moreover, a measure of performance does not give information about the global appearance of the percept, whereas a preference comparison does.

One approach to avoid these issues would be to optimize performance across a diverse range of tasks. Again, our contextual Bayesian optimization framework could be used to adaptively select the most informative tasks. However, since the best encoder may differ depending on the task, this corresponds to a multiobjective optimization problem, which comes with supplementary challenges such as choosing an appropriate trade-off between the different objective optimization problem: the subject balances several aspects of the percept and decides which are the most important, the advantage being that this trade-off between objectives is not arbitrary, but depends only on the subject. This parallels the idea of using preferential optimization in multi-objective engineering design optimization (Lepird et al., 2015). Despite the fact that performance-based optimization comes with many challenges, in the future, it could be useful, for example, to find encoders specifically adapted to some tasks such a face recognition or reading, and the subject could select a specific encoder to handle the task at hand.

5.1.4 Practical preferential Bayesian optimization.

When applying advanced optimization techniques to real-world problems, many factors can impede the success of the algorithm. In our case, three steps had to be carefully considered: the selection of the surrogate model kernel and hyperparameters, the acquisition function, and some aspects of the implementation (in particular the approximate sampling methods). The solution we proposed to these different issues in chapter 3 are general and could be transferred to other problems.

A critical aspect in the development of the method was the use of a model of a virtual patient, mimicking the response of a human in the task. Indeed, in essence, the algorithm cannot be tested on a pre-built dataset since the training set is built online. But it would have been very time-consuming to always rely on human responses over the course of the method development. Instead, we built a simple response model, assuming that the preference is related to distance (Mean Squared Error) between stimulus and percepts. Although this emulator is far from accurately describing the response of real participants, it captures some essential aspects of the task such as dimension, preference feedback, and overall objective smoothness. It proved to be a valuable tool to rapidly assess various algorithms in the exploratory phase of the research.

Recently, there has been considerable interest in automating the process of building machine learning models, with applications for example to hyperparameters (such as learning rate or regularization coefficients) optimization or neural architecture search (Snoek et al., 2012). This is known as Automated Machine Learning (Hutter et al., 2019), and it is clear that such methods would be needed to ease the use of machine learning in a wider range of applications. Bayesian optimization is a powerful tool among the methods that were developed for this purpose. However, automating the process of optimizing a Bayesian optimization method is even more challenging, especially due to the fact that the surrogate model is built online and determines the new training data. Many design options are available, such as the number of initial random sampling steps, the acquisition function, or the surrogate model, which can all affect performance (Lindauer et al., 2019). Probably the most critical aspect is the choice of the surrogate model: its structure, hyperparameters, etc. When using Gaussian processes, Bayesian model selection is usually performed by adapting the kernel hyperparameters online, and recently, some authors generalized this approach to include a nested active model search within the BO loop with promising results (Malkomes and Garnett, 2018; Gardner et al., 2017). Another solution would be, when possible, to use a fast and cheap emulator of the problem at hand, like we did when developing our own method by using a simple model of human response in the task. To our knowledge, this idea of using an emulator for automatically optimizing a Bayesian optimization algorithm for a specific task has never been put into practice, and we believe it has great potential.

5.2 Limitations of our approach

This work gives insight into the limitations of retinal encoders design and Human-in-the-loop optimization for tuning visual prostheses, as well as challenges that may arise when trying this approach on real-life patients and potential avenues to bypass them.

5.2.1 Application to real-life patient

An aspect of prosthetic vision that may limit the applicability of our approach is the variability in percepts evoked for identical stimulation parameters in the same patient. Beyeler et al. (2019b) noted that the same electrode can evoke percepts with very different shapes for two repetitions of the same stimulation protocol. The reason for this is still unknown and this phenomenon has not been extensively studied, but it is clear that if phosphene variability is too high, the mere idea of using a model-based encoder and subjects feedbacks may be pointless.

The encoder design that we proposed is based on a biophysical model of prosthetic vision. However, in practice, there may be some mismatch between the model and reality. We showed that even strong model misspecification did not impede the optimization, even though this impacts the quality of the optimum. Indeed, our strategy is not to learn the model parameters, but instead to optimize these parameters, independently of their physical meaning. This also highlights the importance of loose enough constraints on the encoder: there is an inherent tradeoff between the number of encoder parameters and the encoder flexibility. A more complex encoder is more difficult to optimize given a limited amount of data, but with a tightly constrained encoder, the optimal parameters may correspond to much poorer vision compared to what could be achieved with the same prosthetic device.

Finally, in our experiment, we restricted ourselves to high contrast stimuli. We did not test whether improvement in perceptual experience was transferred to complex stimuli with many contrast levels. To do so, it may be necessary to add more diversity to the stimuli that were presented during the optimization sequence.

5.2.2 Generalization to new generation implants

The *pulse2percept* model has been implemented for a range of implants including PRIMA. However, we choose to focus on Argus II as the model was specifically validated on Argus I and II. Is the approach we proposed applicable to new generation devices such as PRIMA or POLYRETINA (see 1.2), which are made of photovoltaic components? Importantly, the learning algorithm we used is agnostic with respect to the implant or the patient's condition. Only the encoder parameterization is based on detailed knowledge of the implant design and evoked percepts. By inverting the static *pulse2percept* model, we effectively limited axon activation and electrodes crosstalk. A priori, these problems do not arise, or at least arise to a small extent, with technologies such as PRIMA. However, even with recent implants, images are distorted in several ways that could be compensated by appropriate stimulus encoding. This includes for example distortion due to unspecific activation of the ON and OFF pathways (Golden et al., 2019). As a consequence, for our approach to be beneficial to PRIMA users, it would probably be necessary to integrate more details in the perceptual model, such as temporal dynamics, or how retinal ganglion cell types each contribute to image percepts, as well as their respective sensitivity to current.

A method that is fully agnostic concerning the device could be implemented, for example by directly optimizing the encoder matrix *W*. However, in this case, the dimension of the optimization problem would be the number of electrodes times the number of pixels in the input image, making the problem very high dimensional.

5.2.3 Temporal coding

The encoder we used in our experiments did not take into account the temporal dimension, that is the temporal pattern of pulses did not depend on the stimuli. Including the temporal dynamics could for example diminish perceptual fading.

Our preference-based optimization approach could easily generalize to include the temporal dynamic of the percept and tune an encoder including the temporal dimension. In our experiments, the percept was stationary, but this was not a necessary condition. We could for example ask the subject to compare percepts on short periods of time, of the order of a few seconds, so that the subject selects the encoder leading to the most stable percepts.

5.2.4 Non-linearity of the perceptual model

However, a difficult problem to address is the parameterization of the temporal code. Indeed, in our work, it was easy to find a good parameterization for a linear code *W* because of the close-to-linear interactions between electrodes projective fields: $p_{\phi}(a) = \max(|M_{\phi}a|, m)$, where *p* is the subject's percept, *a* the electrodes amplitudes, *m* the maximal percept luminance and M_{ϕ} the matrix whose rows correspond to electrodes projective fields, i.e. how a given electrode contributes to each pixel brightness in the perceptual space, and the eight patient-specific parameters are denoted as ϕ (Beyeler et al., 2019b). With a non-linear mapping from electrodes activity to percepts, one may have to use a non-linear code. Inverting a non-linear perceptual model to find a good non-linear encoder is a problem in itself. Solving this problem may be inevitable for other reasons, for example, to specifically stimulate distinct cell types.

One solution may be simply not to make use of a perceptual model to parameterize the encoder, and this is the strategy we initially pursued. But this makes it very difficult to find a relevant and low-dimensional encoder. In the following section, we describe possible strategies to solve this problem by leveraging domain knowledge.

5.3 Leveraging domain knowledge

Like every complex machine learning task, the ability of the algorithm to perform well is directly determined by the choice of constraints and prior information. We have already discussed the methods we used to select these in chapter 2, but other approaches may also be relevant.

The surrogate Gaussian process model we used in our experiments is defined with a Matérn kernel. Like most kernels used in Bayesian optimization, it is local, i.e. it encodes the assumption that the covariance between pairs of points decreases with their Euclidean distance. However, the average distance between pairs of points selected at random grows exponentially with dimension. So, unless the target data distribution lies in a lower-dimensional manifold, the number of training data required will increase exponentially with dimension (Bengio et al., 2005). This sensitivity of local kernel machines to the curse of dimensionality suggests that our method would need to be adapted to work with higher dimensional encoders.

5.3.1 Informed surrogate model

In Bayesian optimization, prior knowledge about the problem can be included by informing the prior distribution. Among the possibilities, one could consider using an informed Gaussian process prior as a surrogate model, for example by using a non-zero mean that favors some regions of the search space that are a priori promising, reducing the effective size of the space. It has been shown on synthetic examples that the rate of convergence can be significantly improved by choosing an appropriate prior mean (De Ath et al., 2020).

One approach is to parameterize the surrogate model (instead of using a non-parametric Gaussian process). Cox and De Vries (2017) suggested that modeling the objective function with bell, cone, or parabolic shape functions may speed up learning. This idea is based on the intuition that preference functions often have such simple shapes. However, the authors did not show how robust the method was to differences between the parameterization of the objective function and the surrogate model. We tried to use such parameterization in pilot experiments. We modeled the preference using a Gaussian function, with an unknown mean and covariance matrix. We put a Gaussian prior on the mean and a Wishart prior on the inverse covariance matrix (the Wishart distribution is the conjugate prior of the inverse covariance matrix of a multivariate normal distribution). It did not perform as well as the method based on Gaussian processes. This may be because the parameterized surrogate lacks flexibility, or even because the single-peak assumption does not hold.

Another complementary approach is to use previous data and domain knowledge to constrain the search space as much as possible. Perrone et al. (2019) proposed a very general procedure to learn an ellipsoid search space on a set of previous experiments. Concretely, we could learn the distribution of optimal encoder parameters over the population of patients and compute the smallest ellipsoid that encloses these parameters. Then, this ellipsoid would define the search space for the next patient. Unless the optimum for the current patient is an outlier with respect to the rest of the population, this approach should improve the speed of convergence. This is related to what we did in our experiment, as we used data from previous studies to constrain the search space.

5.3.2 Structured kernels

Structure discovery

To keep the advantages of Gaussian process models in Bayesian optimization in high dimensions without an exponential increase in available data, one possibility is to use kernels encoding non-local interactions, called non-local kernels. The advantage of using non-local kernels is that, since they encode dependency between distant points, they allow extrapolating away from training data. A simple method to design such covariance functions is the use of additive kernels (Duvenaud, 2014), which reflect an additive structure in the underlying function. Additivity means that the function can be written as the sum over functions of subsets of variables. Although restricting, this assumption holds true for many real-world applications (Duvenaud, 2014). This structure can be automatically discovered along the optimization process (Gardner et al., 2017; Wang et al., 2017b; Mutný and Krause, 2018; Rolland et al., 2018). A drawback, however, is that in general, it is not possible to know in advance whether such structure exists in the objective. Another strategy is to try to discover a lower-dimensional structure in the data-generating function (Zhang et al., 2019). The method works by linearly projecting the kernelized inputs to a lower-dimensional subspace using Sliced Inversion Regression. A Gaussian process model is learned on these transformed inputs and acquisition is performed in the standard way. This method showed good performance in high-dimensional synthetic problems. If in real-life applications the model-based encoder parameterization appears too constraining, both approaches could be investigated to loose constraints on the encoder.

Sim-to-real transfer for kernel engineering

Another approach to designing the kernel without having to hand-tune it is the use of a form of transfer learning called sim-to-real transfer. Transfer learning is the training of a model that attempts to reuse knowledge of the source tasks for the target task. Sim-to-real transfer is a form of transfer learning that uses simulated data. It has been primarily used in teaching robots complex motor skills (Golemo et al., 2018). The method essentially works by making robots learn motor policies in a simulated environment. Because of the mismatch between simulation and reality, a crucial aspect is to learn policies that transfer to the real world without preventing the ability of the robot to adapt afterward to real-world tasks. A method to inform Bayesian optimization with simulations was proposed in a series of publications by Antonova et al. (2017) and Rai et al. (2019) to optimize a robot's locomotion controller. The idea is to learn an informed kernel using simulations. A deep neural network model ϕ is trained to predict the performance of the robot as a function of the control parameters x using simulated trajectories: score(x) = $\phi(x, w)$. w denotes the parameters of the deep network that are learned using simulated data.

This neural network is then used to construct an informed deep kernel $k_{\text{deep}}(x, x'|w) = k(\phi(x, w), \phi(x', w))$, where *k* is a standard local kernel, like the squared exponential kernel.

In a sense, this performance metric is used to warp the input space, so that promising parameters are closer, and far away from unpromising ones. So here, the use of deep kernels is closely related to the more general problem of metric learning (Bellet et al., 2013). The advantage of using Gaussian processes with deep kernels is that it leverages the power of deep learning to learn hidden structures while keeping the flexibility of kernel methods.

Another, more flexible method to learn deep kernels was introduced by Wilson et al. (2016), which allows to simultaneously learn a Gaussian process and its deep kernel (whereas the former method trains the deep network and the Gaussian process independently).

The sim-to-real transfer strategy could be applied to the problem of tuning retina implant, using prosthetic vision models like *pulse2percept* (Beyeler et al., 2017a). However, to do so, we would need a metric that relates the predicted percepts to perceptual quality, which in itself is a challenge, given that standard image quality metrics (Simoncelli et al., 2004) may not be relevant for prosthetic vision. Our observations indeed suggest that Mean-Squared Error (MSE) between percepts and stimuli is not a relevant metric to predict prosthetic percept quality (chapter 2). One interpretation of this result is that for such low-quality vision, perception of high-level features (such as the edges of the letters) is much more important than a pixel-by-pixel agreement between input and percept.

Despite this problem of defining a relevant metric for perceptual quality, we investigated the use of sim-to-real transfer in preliminary work. For a given set of true parameters ϕ^* , we simulated percepts for random encoders and trained a Gaussian process with a deep kernel to predict the MSE between the percept and the input image. The idea is then to use this informed kernel in Bayesian optimization of the retinal implant encoder in simulations. The core idea behind this strategy is that, although there is variability between subjects, the structure of the latent value function for pairs of parameters is approximately the same from one subject to another. This is a more evolved version of what we did in the experiments to learn and transfer hyperparameters between subjects. Although we believe this approach is promising, this work is still at an embryonic stage. A significant challenge is to find a neural network architecture able to efficiently learn the mapping and to generalize from one patient to another. Indeed, in the transfer learning strategy we used in practice, we noted that the Squared Exponential kernel with Automatic Relevance Determination (SE-ARD) was better at predicting the subject's responses on the training set, but not on the test set. This over-fitting is probably explained by the fact that the ARD kernel has more degrees of freedom (9 instead of 2 for the Matérn 5/2 kernel). One way to circumvent this problem would be to train the deep kernel by using simulations from a wide range of perceptual models. Of course, provided clinicians gather enough data from subjects, the same strategy could be applied by replacing simulated data with real data.

The challenge of high-dimensional acquisition

Another issue with high-dimensional Bayesian optimization is the possibly prohibitive cost of global optimization of the acquisition function at each iteration. This is especially a problem in HILO since, contrary to standard optimization where the cost is solely incurred by function evaluation, in that case, the cost is the total time spent by the subject in the optimization procedure, which includes the acquisition. One proposal is to optimize only a subset of the variables at each iteration (Li et al., 2017), which showed good performance on synthetic problems while speeding up acquisition. Another possibility is to perform asynchronous parallel acquisition, where several next inputs are computed in parallel and a query is performed whenever an acquisition terminates (at a given time, different acquisition computations do not necessarily work with the same amount of data, hence the name asynchronous, see e.g. Kandasamy et al. (2018)).

5.4 Getting more from participants

As we showed in our experiment, the type of feedback from patients is paramount in the success of the encoder optimization procedure. We used pairwise comparisons for the reasons explained above, but other types of feedbacks may be more efficient.

5.4.1 Ties inclusion

An obvious limitation of the algorithm we used is that it does not take ties into account. That is, when the subject finds percepts of equal quality, he/she is forced to make a choice between the two. Indeed, subjects reported that they sometimes had to answer at random. We did not quantify the fractions of queries when this occurred, but allowing for ties may decrease the model's uncertainty and increase optimization robustness. One reason that made us choose not to include ties in our experiments is that participants, destabilized by the appearance of the percepts, may prefer not to judge instead of making efforts to find differences between the two alternatives. In a way, we wanted the participants to learn to judge these very unusual stimuli.

Dewancker et al. (2017) used the generalized Bradley-Terry model (Rao and Kupper, 1967) to include ties in PBO. However, they used a mean-field approximation strategy for inference that implies that the posterior latent function values are independent. Recently, Nguyen et al. (2020) proposed a more efficient approach based on the multinomial logit model.

5.4.2 From pairwise comparisons to top-k ranking

Another method that may be efficient to get more information from patients would be to ask them to select the best configuration among a limited set of possibilities. Of course, a trade-off would have to be found between information gain and cognitive demand, especially in terms of memory. Indeed, for implant users, parameters have to be evaluated sequentially, which implies remembering the experience they had with previous configurations.

The scenario of "multi-duels" has been investigated by Sui and Burdick (2017) and Siivola et al. (2020), and we proposed an improved implementation as well as a more efficient acquisition function in 3.5. Sui and Burdick (2017) and Siivola et al. (2020) showed that this type of feedback greatly fastens optimization compared to pairwise comparisons. Intuitively, this stems from the fact that each feedback is more informative since it corresponds to simultaneously performing multiple pairwise comparisons (see section 3.5). Nguyen et al. (2020) later generalized PBO beyond pairwise preference judgments to include top-k rankings (that is, a ranking of the top-k inputs in a finite set).

5.4.3 Continuous tuning

Imagine that we continuously and simultaneously tune the parameters of the implant to find the preferred configuration. This continuous tuning could greatly increase the ease and speed of implants tuning. Mikkola et al. (2020) devised a method they called Projective Preferential Bayesian Optimization, where the user tunes a single variable, which is a linear combination of all the parameters. This way, the parameters of the system change in a pre-determined direction. Once the user has found the best configuration along this direction, a new one is proposed. Another related approach is "Gibbs sampling with people" (Harrison et al., 2020), where parameters are continuously tuned one after the other.

However, for these approaches to be beneficial compared to a discontinuous exploration of the parameters space, it requires that the time it takes for a patient to evaluate a given configuration is short. Indeed, this method is much less sample-efficient. In our experiments, the limiting factor was the time it took to compute the encoder corresponding to the new candidate parameters. Indeed, computing this encoder required estimating each electrode projective field using the *pulse2percept* model, which took a few seconds. As a consequence, continuous tuning seemed inappropriate. However, this could be solved in the future by substituting fast emulators to complex biophysical models. One approach for example would be to train a fast statistical model to predict M_{ϕ} for any ϕ , using a dataset generated using *pulse2percept*.

5.4.4 Everyday tuning

Automatic optimization opens the door to out-of-the-lab implant tuning, with the patients adjusting the stimulation protocol on their own. With preference-based optimization, patients could continuously adjust their preference as they learn to use their prosthetic vision and as their visual system changes due to disease progression. The patient could simply select a learning mode whereby new candidate encoders are selected by the preferential optimization algorithm and assessed by the user. In principle, this would also be possible with performancebased optimization, provided the patient is equipped with a system that runs a battery of tests and measures the patient's performance.

Such autonomous implant tuning would be especially useful in case of implant malfunction, with some electrodes becoming inactive, This possibility could be investigated in healthy subjects by equipping them with a virtual reality device so that they see the world in simulated prosthetic vision and adjust the encoder in daily tasks.

5.5 Trends in Bayesian optimization

Bayesian optimization is a very active research area, with a recent surge of interest driven by its performance in applications such as tuning of models hyperparameters (Snoek et al., 2012; Turner et al., 2021). Innovation in this field is driven in two main ways: improvement in surrogate modeling, and in acquisitions.

Improved surrogate modeling

On the first aspect of surrogate modeling, a major trend, that I previously mentioned, is to scale to high dimensions without losing the good uncertainty quantification abilities and flexibility of Gaussian processes. Another, related research direction, is to use progress in learning representations and transfer learning, as illustrated for example by the use of deep kernels. In most applications, simplifying assumptions such as stationary and homoscedasticity are made, which may not hold in practice. Methods to handle these difficulties have been developed, including in particular warped Gaussian processes (Lázaro-Gredilla, 2012), which learn a non-linear transformation of the Gaussian process, input warping (Snoek et al., 2014), based on learning a non-linear transformation of the input space, and deep Gaussian processes (Damianou and Lawrence, 2013), which correspond to the successive composition of Gaussian processes.

However, as I already mentioned, more complex models require more data and constraints to avoid overfitting. Hence, given the limited amount of data that is available in a typical BO task, methods to include prior knowledge, such as transfer learning, are essential. For example, the use of deep kernel surrogates trained on related problems showed excellent performance on hyperparameters optimization tasks (Wistuba and Grabocka, 2021).

Improved acquisition

Acquisition rules are mostly heuristics, and given the problem complexity, the first regret bounds were only published quite recently (Srinivas et al., 2010). For complex scenarios such as binary feedbacks, no such proofs are available. It is clear that a deeper theoretical understanding will be beneficial to design acquisition rules. Among the recent trends in designing acquisition strategies, the first one is non-myopic BO (González et al., 2016; Wu and Frazier, 2019; Jiang et al., 2019), in which at each iteration, the selection takes into account the fact that other iterations will follow. It is possible to show on simple examples that non-myopic BO tends to explore more aggressively, and shows better performance (Wu and Frazier, 2019). The main challenge to address is the computational cost of non-myopic heuristics.

Another trend is in designing acquisition methods that take into account the uncertainty about the surrogate model, and in particular, the kernel's hyperparameters. For example, a version of Thomspon sampling was proposed where the hyperparameters are sampled from their posterior distribution, and a point is then sampled from the posterior distribution over the maximum, conditionally on the hyperparameters (De Palma et al., 2019).

Finally, a promising approach is, instead of designing acquisition rules, to 'learn to acquire'. By meta-learning acquisition functions on a set of tasks, one could make the algorithm particularly efficient on related tasks. This strategy was introduced by Volpp et al. (2019), who trained a 'neural acquisition function': a neural network taking the posterior variance and posterior mean of a GP surrogate as input, and outputting a sample for the next iteration. The neural network was trained using reinforcement learning. This approach showed good performance in numerical experiments.

5.6 Conclusion

Progress in the field of machine learning is and will continue to be paramount in improving the performance of brain-machine interfaces. It remains to be proven if the Bayesian approach we described here will be efficient in real-world patients. We hope that this thesis will contribute to lightening a path where models of prosthetic vision, clinical knowledge, and patients' experience will be articulated to bring vision back to those who lost it.

A

Bayesian analysis of preference data

In order to compare preference between encoders, we used the following Bayesian analysis method adapted from Kass and Raftery (1995) (see also chapter 5 of Robert (2006)): For two algorithms A_1 and A_2 used to derive encoders, we have the two hypotheses H_1 (resp. H_2): encoders found using A_1 (resp. A_2) are preferred to encoders found using A_2 (resp. A_1). For a participant *i*, the responses in a series of preference comparisons between two encoders follows a binomial distribution $y_i = \mathcal{B}(n_i, p_i)$, where n_i is the number of comparisons and y_i is the number of comparisons for which the participant reported that A_2 is better than A_1 . We put the conjugate beta prior $p_i \sim Beta\left(\frac{\xi}{\omega}, \frac{1-\xi}{\omega}\right)$, where $\xi \in [0, 1]$ and $\omega \in \mathbb{R}^*_+$. Note that $\xi = \mathbb{E}(p_i)$. We assume that ω is constant, and ξ follows a uniform distribution. Under H_1 , the model is $\xi \sim \mathcal{U}([0, 0.5])$, whereas under H_2 , $\xi \sim \mathcal{U}([0.5, 1])$.

The Bayes factor B_{21} is the ratio between the likelihoods of the two models: $B_{21} = \frac{p(\mathscr{D}|H_2)}{p(\mathscr{D}|H_1)}$. It gives an estimate of how strongly the data support H_2 over H_1 .

We have:

$$p(\mathscr{D}|H_i) = \int_0^1 p(\mathscr{D}|\xi) p(\xi|H_i) d\xi.$$
(A.1)

So for H_1 , by noting *G* the number of participants:

$$p(\mathscr{D}|H_1) = 2 \int_0^{0.5} \prod_{i=1}^G \int_0^1 p_i^{y_i} (1-p_i)^{n_i - y_i} p_i^{\alpha - 1} (1-p_i)^{\beta - 1} dp_i \times \left[\frac{\Gamma(\frac{1}{\omega})}{\Gamma(\frac{\xi}{\omega})\Gamma(\frac{1-\xi}{\omega})} \right] d\xi, \qquad (A.2)$$

where $\alpha = \frac{\xi}{\omega}$ and $\beta = \frac{1-\xi}{\omega}$. So:

$$p(\mathscr{D}|H_1) = 2 \int_0^{0.5} \prod_{i=1}^G \Gamma\left(y_i + \frac{\xi}{\omega}\right) \frac{\Gamma(n_i - y_i + \frac{1 - \xi}{\omega})}{\Gamma(n_i + \omega^{-1})} \times \left[\frac{\Gamma(\frac{1}{\omega})}{\Gamma(\frac{\xi}{\omega})\Gamma(\frac{1 - \xi}{\omega})]}\right] d\xi$$
$$= 2 \int_0^{0.5} \prod_{i=1}^G \left[\prod_{j=1}^{y_i} \left(j - 1 + \frac{\xi}{\omega}\right) \prod_{j=1}^{n_i - y_i} \left(j - 1 + \frac{1 - \xi}{\omega}\right) \left(\prod_{j=1}^{n_i} \left(j - 1 + \omega^{-1}\right)\right)^{-1}\right].$$
(A.3)

So the integrand in the likelihood is simply a polynomial function of ξ that can easily be integrated. Similarly, for the alternative hypothesis H_2 :

$$p(\mathscr{D}|H_2) = 2 \int_{0.5}^{1} \prod_{i=1}^{G} \left[\prod_{j=1}^{y_i} \left(j - 1 + \frac{\xi}{\omega} \right) \prod_{j=1}^{n_i - y_i} \left(j - 1 + \frac{1 - \xi}{\omega} \right) \left(\prod_{j=1}^{n_i} \left(j - 1 + \omega^{-1} \right) \right)^{-1} \right].$$
(A.4)

For each comparison between algorithms, ω was estimated empirically, by noting that: $\mathbb{V}(p_i) = \frac{\omega\xi(1-\xi)}{1+\omega}$. So we computed the empirical mean of the p_i , $\hat{\xi} = \langle p_i \rangle$ and the empirical variance $\hat{V} = \langle p_i - \hat{\xi} \rangle$, and solved for $\hat{V} = \frac{\omega\xi(1-\xi)}{1+\omega}$. We then evaluated the integral numerically. An advantage of this analysis method is that it models the noise in the preference measurement data.

B

A standard result gives the distribution of a subset of dimensions y_A of a vector y having a multivariate Gaussian distribution, conditioned on the values of another subset y_B . A proof can be found in section 2.3.1 of Bishop (2006). If

$$\boldsymbol{y} = \begin{bmatrix} \boldsymbol{y}_A \\ \boldsymbol{y}_B \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \boldsymbol{\mu}_A \\ \boldsymbol{\mu}_B \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{AA} & \boldsymbol{\Sigma}_{AB} \\ \boldsymbol{\Sigma}_{BA} & \boldsymbol{\Sigma}_{BB} \end{bmatrix} \right), \quad (B.1)$$

then

$$\boldsymbol{y}_A \mid \boldsymbol{y}_B \sim \mathcal{N} \left(\boldsymbol{\mu}_A + \boldsymbol{\Sigma}_{AB} \boldsymbol{\Sigma}_{BB}^{-1} \left(\boldsymbol{y}_B - \boldsymbol{\mu}_B \right), \boldsymbol{\Sigma}_{AA} - \boldsymbol{\Sigma}_{AB} \boldsymbol{\Sigma}_{BB}^{-1} \boldsymbol{\Sigma}_{BA} \right).$$
 (B.2)

This result can be used in the context of Gaussian process regression. With a GP prior $\mathcal{GP}(m(\cdot), k(\cdot, \cdot))$ on f. If we denote (X, y) the observations and X^* the test points, the covariance matrices are given by:

$$\begin{split} \boldsymbol{\Sigma}_{AA} &= k\left(\boldsymbol{X}^{*}, \boldsymbol{X}^{*}\right), \\ \boldsymbol{\Sigma}_{AB} &= k\left(\boldsymbol{X}^{*}, \boldsymbol{X}\right), \\ \boldsymbol{\Sigma}_{BA} &= k\left(\boldsymbol{X}, \boldsymbol{X}^{*}\right), \\ \boldsymbol{\Sigma}_{BB} &= k\left(\boldsymbol{X}, \boldsymbol{X}\right), \end{split}$$

and the means are given by:

$$\mu_A = m(X^*),$$
$$\mu_B = m(X).$$

As a consequence, the predictive distribution of f at the test points conditioned on the observations is:

$$f(X^*) \mid (\mu(X^*), \Sigma(X^*, X^*)),$$
 (B.3)

with:

$$\mu(\mathbf{X}^*) = m(\mathbf{X}^*) + k(\mathbf{X}^*, \mathbf{X})k(\mathbf{X}^*, \mathbf{X}^*)^{-1}(\mathbf{y} - m(\mathbf{X})),$$

$$\Sigma(\mathbf{X}^*, \mathbf{X}^*) = k(\mathbf{X}^*, \mathbf{X}^*) - k(\mathbf{X}^*, \mathbf{X})k(\mathbf{X}, \mathbf{X})^{-1}k(\mathbf{X}, \mathbf{X}^*).$$
(B.4)

In the following definitions, ρ is the lengthscale of the kernels (positive constant), and σ are scaling parameters.

The Matérn kernels is defined as:

$$k_{\nu}(\boldsymbol{x},\boldsymbol{x'}) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{\|\boldsymbol{x} - \boldsymbol{x'}\|}{\rho}\right)^{\nu} K_{\nu} \left(\sqrt{2\nu} \frac{\|\boldsymbol{x} - \boldsymbol{x'}\|}{\rho}\right), \quad (C.1)$$

where Γ is the gamma function, K_{ν} is the modified Bessel function of the second kind and ν is a positive parameter. For example, the Matérn 5/2 kernel is defined as:

$$k_{\nu=5/2}(\boldsymbol{x}, \boldsymbol{x}') = \sigma^2 \left(1 + \frac{\sqrt{5} \|\boldsymbol{x} - \boldsymbol{x}'\|}{\rho} + \frac{5 \|\boldsymbol{x} - \boldsymbol{x}'\|^2}{3\rho^2} \right) \exp\left(-\frac{\sqrt{5} \|\boldsymbol{x} - \boldsymbol{x}'\|}{\rho}\right).$$
(C.2)

The Squared-Exponential kernel is defined as:

$$k_{SE}(\boldsymbol{x}, \boldsymbol{x}') = \sigma^2 \exp\left(-\frac{\parallel (\boldsymbol{x} - \boldsymbol{x}') \parallel^2}{2\rho}\right).$$
(C.3)

For stationary local kernels such as the Matérn and the Squared-Exponential kernels, it is possible to specify a different lengthscale for each input dimension. This is known as Automatic Relevance Determination (ARD). For example, the SE-ARD kernel is defined as:

$$k_{SE}(\boldsymbol{x}, \boldsymbol{x'}) = \sigma^2 \exp\left(-\frac{1}{2}(\boldsymbol{x} - \boldsymbol{x'})^\top \boldsymbol{\Lambda}(\boldsymbol{x} - \boldsymbol{x'})\right), \qquad (C.4)$$

where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_D)$. *D* is the dimension of the input space and the λ_i are positive constants.

The linear kernel is defined as:

$$k_{\text{lin}}(\boldsymbol{x}, \boldsymbol{x}') = \sigma_b^2 + \sigma^2 (\boldsymbol{x} - \boldsymbol{c})^\top (\boldsymbol{x}' - \boldsymbol{c}) , \qquad (C.5)$$

where the offset *c* determines the coordinate of the intercepts of all samples from the posterior, and σ_b determines how far from 0 the function will be at zero(Duvenaud, 2014).

And the periodic kernel is defined as:

$$k_{\text{per}}(\boldsymbol{x}, \boldsymbol{x'}) = \sigma^2 \exp\left(-\frac{2\sin^2\left(\pi \left\|\boldsymbol{x} - \boldsymbol{x'}\right\|/p\right)}{\rho}\right),\tag{C.6}$$

where *p* determines the periodicity of the function.

D

Gradient of the variance of the predictive class distribution for GP classification

We start by computing the variance of the predictive class distribution for a Gaussian process classification model. By definition:

$$\mathbb{V}\left[\Phi(f(\mathbf{x}))\right] = \int \left(\Phi(f(\mathbf{x})) - \mathbb{E}\left[\Phi(f(\mathbf{x}))\right]\right)^2 p(f \mid \mathcal{D}_t, \mathbf{x}) df$$
(D.1)

As a consequence:

$$\mathbb{V}\left[\Phi(f(\mathbf{x}))\right] = \int \Phi(f(\mathbf{x}))^2 p(f(\mathbf{x}) \mid \mathscr{D}_t, \mathbf{x}) df - \mathbb{E}\left[\Phi(f(\mathbf{x}))\right]^2 \\
= \int \Phi\left(\sigma_f(\mathbf{x})f(\mathbf{x}) + \mu_f(\mathbf{x})\right)^2 \phi\left(f(\mathbf{x})\right) df - \mathbb{E}\left[\Phi\left(f(\mathbf{x})\right)\right]^2,$$
(D.2)

where ϕ is the probability density function of $\mathcal{N}(0,1)$, μ_f and σ_f^2 are the mean and variance of the Gaussian approximation to the posterior $p(f|\mathcal{D}_t)$.

Making use of the following formula (Patel and Read, 1996):

$$\int_{-\infty}^{\infty} \Phi(a+bt)^2 \phi(t) dt = \Phi\left(\frac{a}{\sqrt{1+b^2}}\right) - 2T\left(\frac{a}{\sqrt{1+b^2}}, \frac{1}{\sqrt{1+2b^2}}\right),$$
 (D.3)

with the Owen's T function defined as:

$$T(h,a) = \frac{1}{2\pi} \int_0^a \frac{e^{-\frac{1}{2}h^2(1+t^2)}}{1+t^2} dt,$$
 (D.4)

equation D.2 can be written as:

$$\mathbb{V}_{f}\left[\Phi\left(f(\boldsymbol{x})\right)\right] = \mathbb{E}_{f}\left[\Phi\left(f(\boldsymbol{x})\right)\right] - 2T\left(\frac{\mu_{f}(\boldsymbol{x})}{\sqrt{1 + \sigma_{f}^{2}(\boldsymbol{x})}}, \frac{1}{\sqrt{1 + 2\sigma_{f}^{2}(\boldsymbol{x})}}\right) - \mathbb{E}_{f}\left[\Phi\left(f(\boldsymbol{x})\right)\right]^{2}.$$
 (D.5)

The derivatives of the Owen's T function are:

$$\frac{\partial T}{\partial h} = -\frac{1}{2\pi} \int_0^a h e^{-\frac{1}{2}h^2(1+x^2)} dx = -\frac{1}{2\pi} e^{-\frac{h^2}{2}} \int_0^{ha} e^{-\frac{t^2}{2}} dt = -\frac{e^{-\frac{h^2}{2}}}{\sqrt{2\pi}} \left(\Phi(ha) - \frac{1}{2}\right), \quad (D.6)$$

$$\frac{\partial T}{\partial a} = \frac{1}{2\pi} \frac{e^{-\frac{1}{2}h^2(1+a^2)}}{1+a^2}.$$
 (D.7)

Moreover, we have (see appendix I):

$$\mathbb{E}_{f}\left[\Phi\left(f(\boldsymbol{x})\right)\right] = \Phi\left(\frac{\mu_{f}(\boldsymbol{x})}{\sqrt{1 + \sigma_{f}^{2}(\boldsymbol{x})}}\right).$$
(D.8)

So computing the gradient of $\mathbb{V}_f[\Phi(f(\mathbf{x}))|\mathcal{D}_t]$ is straightforward.

E

The knowledge gradient is defined as:

$$\mathrm{KG}(\mathbf{x}) = \mathbb{E}_{c \sim p(c|\mathscr{D})}(\mu_{n+1}^{\star} - \mu_n^{\star}|\mathscr{D}, \mathbf{x}_{n+1} = \mathbf{x}), \tag{E.1}$$

where μ_{n+1}^{\star} is the maximum of the posterior mean after observing $(\mathcal{D}, (\mathbf{x}, c))$, and μ_n^{\star} is the maximum of the posterior mean after observing $\mathcal{D} = (\mathbf{X}, c_{1\cdots,n})$.

With binary outputs, the knowledge gradient can be expressed as:

$$KG(\mathbf{x}) = \mu_c(\mathbf{x})(\mu_{1,n+1}^{\star} - \mu_n^{\star}) + (1 - \mu_c(\mathbf{x}))(\mu_{0,n+1}^{\star} - \mu_n^{\star}),$$
(E.2)

where $\mu_c(\mathbf{x}) = \mathbb{E}_f[\Phi(f(\mathbf{x}))|\mathcal{D}]$, and $\mu_{1,n+1}^{\star}$ (resp. $\mu_{0,n+1}^{\star}$) is the maximum of the posterior mean after observing ($\mathcal{D}, (\mathbf{x}, 1)$) (resp. ($\mathcal{D}, (\mathbf{x}, 0)$)). That is:

$$\mu_{c,n+1}^{\star} = \max_{\mathbf{x}' \in \mathcal{X}} \mathbb{E}_f[f(\mathbf{x}') | \mathcal{D} \cup (\mathbf{x}, c)].$$
(E.3)

The gradient of the knowledge-gradient is given by:

$$\nabla \mathrm{KG}(\mathbf{x}) = \nabla \mu_c(\mathbf{x})(\mu_{1,n+1}^{\star} - \mu_{0,n+1}^{\star}) + \mu_c(\mathbf{x})\nabla \mu_{1,n+1}^{\star} + (1 - \mu_c(\mathbf{x}))\nabla \mu_{0,n+1}^{\star}.$$
 (E.4)

To compute the gradients of $\mu_{1,n+1}^*$ and $\mu_{0,n+1}^*$, inspired by Wu and Frazier (2016), we use the envelope theorem (Milgrom and Segal, 2002), which states that, under sufficient regularity conditions, the gradient with respect to x of a maximum of a collection of functions of x is given simply by first finding the maximum x^* in this collection, and then differentiating this single function with respect to x, keeping x^* fixed. Here, we have an infinite collection of functions $\mathbb{E}_f[f(x')|\mathcal{D} \cup (x,c)]$ indexed by x'.

Theorem 1 Corollary 4 of the envelope theorem in Milgrom and Segal (2002) Let X denote the choice set and t be a parameter in a [0, 1] (the theorem generalizes to normed vector spaces). Let $f : X \times [0, 1] \rightarrow \mathbb{R}$ be an objective function parameterized by t. We define:

$$V(t) = \sup_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, t), \tag{E.5}$$

$$X^{\star}(t) = \{ x \in \mathcal{X}, f(x, t) = V(t) \}.$$
 (E.6)

Suppose that X is a nonempty compact space, f(x, t) is upper semicontinuous in x, and $\frac{\partial f}{\partial t}(x, t)$ is continuous in (x, t). Then:

$$\forall t \in [0,1), V'(t+) = \max_{\mathbf{x} \in \mathcal{X}^*(t)} \frac{\partial f}{\partial t}(\mathbf{x}, t), \tag{E.7}$$

153

Appendix E. Gradient of the knowledge gradient

$$\forall t \in (0,1], V'(t-) = \min_{\mathbf{x} \in \mathcal{X}^{\star}(t)} \frac{\partial f}{\partial t}(\mathbf{x}, t).$$
(E.8)

V is differentiable at any $t \in (0,1)$ if and only if $\left\{\frac{\partial f}{\partial t}(\mathbf{x},t) | \mathbf{x} \in X^{\star}(t)\right\}$ is a singleton, and in that case $\forall \mathbf{x} \in X^{\star}(t), V'(t) = \frac{\partial f}{\partial t}(\mathbf{x},t).$

As a consequence, by writing $x^* = \underset{x' \in \mathcal{X}}{\arg \max} \mathbb{E}_f[f(x') | \mathcal{D} \cup (x, c)]$, we get:

$$\nabla_{\mathbf{x}}\mu_{c,n+1}^{\star} = \nabla_{\mathbf{x}}\mathbb{E}_{f}[f(\mathbf{x}^{\star})|\mathscr{D}\cup(\mathbf{x},c)].$$
(E.9)

From the Laplace approximation, we have:

$$\mathbb{E}_{f}[f(\boldsymbol{x}^{\star})|\mathscr{D}\cup(\boldsymbol{x},\boldsymbol{c})] = \boldsymbol{k}^{\top}\nabla_{\boldsymbol{y}}\log p(\boldsymbol{c}|\boldsymbol{y}), \tag{E.10}$$

were *y* corresponds to the inferred latent values of the training data, $c = c_{1\dots,n+1}$, and $k = k((X, x), x^*)$.

So that:

$$\nabla_{\mathbf{x}} \mathbb{E}[f(\mathbf{x}^{\star})|\mathscr{D} \cup (\mathbf{x}, c)] = (\nabla_{\mathbf{x}} \mathbf{k}^{\top}) \nabla_{\mathbf{y}} \log p(\mathbf{c}|\mathbf{y}) + \mathbf{k}^{\top} \nabla_{\mathbf{x}} \nabla_{\mathbf{y}} \log p(\mathbf{c}|\mathbf{y})$$
$$= (\nabla_{\mathbf{x}} \mathbf{k}^{\top}) \nabla_{\mathbf{y}} \log p(\mathbf{c}|\mathbf{y}) + \mathbf{k}^{\top} (\nabla_{\mathbf{y}}^{2} \log p(\mathbf{c}|\mathbf{y})) \nabla_{\mathbf{x}} \mathbf{y}$$
$$= (\nabla_{\mathbf{x}} \mathbf{k}^{\top}) \nabla_{\mathbf{y}} \log p(\mathbf{c}|\mathbf{y}) - \mathbf{k}^{\top} \mathbf{W} \nabla_{\mathbf{x}} \mathbf{y},$$
(E.11)

where $W = -\nabla_y^2 \log p(c|y)$.

From the Laplace approximation, we have $y = K \nabla_y \log p(c|y)$, where K = k((X, x), (X, x)). By differentiating this self-consistent equation,

$$\nabla_{\mathbf{x}} \mathbf{y} = (\nabla_{\mathbf{x}} \mathbf{K}) \nabla_{\mathbf{y}} \log p(\mathbf{c}|\mathbf{y}) + \mathbf{K} (\nabla_{\mathbf{y}}^2 \log p(\mathbf{c}|\mathbf{y})) \nabla_{\mathbf{x}} \mathbf{y}$$

= $(\nabla_{\mathbf{x}} \mathbf{K}) \nabla_{\mathbf{y}} \log p(\mathbf{c}|\mathbf{y}) - \mathbf{K} \mathbf{W} \nabla_{\mathbf{x}} \mathbf{y}.$ (E.12)

By rearranging:

$$\nabla_{\mathbf{x}} \mathbf{y} = (\mathbf{I} + \mathbf{K} \mathbf{W})^{-1} \nabla_{\mathbf{x}} \mathbf{K} \nabla_{\mathbf{y}} \log(\mathbf{c} | \mathbf{y}).$$
(E.13)

F

F.1 Preference kernel

The mean function m_{pref} and covariance function k_{pref} of the GP prior on the preference function f can be computed from the mean function m and covariance function k of the GP prior on the value function g as follows (Houlsby et al., 2012):

$$k_{\text{pref}}\left(\left[\boldsymbol{x}_{i}, \boldsymbol{x}_{k}\right], \left[\boldsymbol{x}_{j}, \boldsymbol{x}_{l}\right]\right) = \text{Cov}\left[f(\boldsymbol{x}_{i}, \boldsymbol{x}_{k}), f(\boldsymbol{x}_{j}, \boldsymbol{x}_{l})\right]$$
$$+ \text{Cov}\left[g(\boldsymbol{x}_{i}) - g(\boldsymbol{x}_{k}), g(\boldsymbol{x}_{j}) - g(\boldsymbol{x}_{l})\right]$$
$$= \text{Cov}\left[g(\boldsymbol{x}_{i}), g(\boldsymbol{x}_{j})\right] - \text{Cov}\left[g(\boldsymbol{x}_{i}), g(\boldsymbol{x}_{l})\right] - \text{Cov}\left[g(\boldsymbol{x}_{k}), g(\boldsymbol{x}_{j})\right]$$
$$+ \text{Cov}\left[g(\boldsymbol{x}_{k}), g(\boldsymbol{x}_{l})\right]$$
$$= k(\boldsymbol{x}_{i}, \boldsymbol{x}_{j}) + k(\boldsymbol{x}_{k}, \boldsymbol{x}_{l}) - k(\boldsymbol{x}_{i}, \boldsymbol{x}_{l}) - k(\boldsymbol{x}_{k}, \boldsymbol{x}_{j}),$$

and

$$m_{\text{pref}}(\mathbf{x}_i, \mathbf{x}_j) = \mathbb{E}[f(\mathbf{x}_i, \mathbf{x}_j)] = \mathbb{E}[g(\mathbf{x}_i) - g(\mathbf{x}_j)] = m(\mathbf{x}_i) - m(\mathbf{x}_j).$$

By definition, k_{pref} generates valid covariance matrices. The non-stationarity of the preference kernel can easily be verified:

$$k_{\text{pref}} ((\mathbf{x}_{i}, \mathbf{x}_{j}) + (\mathbf{x}_{s}, \mathbf{x}_{s}'), (\mathbf{x}_{k}, \mathbf{x}_{l}) + (\mathbf{x}_{s}, \mathbf{x}_{s}'))$$

$$= k(\mathbf{x}_{i} + \mathbf{x}_{s}, \mathbf{x}_{k} + \mathbf{x}_{s}) + k(\mathbf{x}_{j} + \mathbf{x}_{s}', \mathbf{x}_{l} + \mathbf{x}_{s}') - k(\mathbf{x}_{i} + \mathbf{x}_{s}, \mathbf{x}_{l} + \mathbf{x}_{s}') - k(\mathbf{x}_{i} + \mathbf{x}_{s}, \mathbf{x}_{l} + \mathbf{x}_{s}') - k(\mathbf{x}_{j} + \mathbf{x}_{s}', \mathbf{x}_{k} + \mathbf{x}_{s})$$

$$= k(\mathbf{x}_{i}, \mathbf{x}_{k}) + k(\mathbf{x}_{j}, \mathbf{x}_{l}) - k(\mathbf{x}_{i} + \mathbf{x}_{s}, \mathbf{x}_{l} + \mathbf{x}_{s}') - k(\mathbf{x}_{j} + \mathbf{x}_{s}', \mathbf{x}_{k} + \mathbf{x}_{s})$$

$$\neq k_{\text{pref}} ((\mathbf{x}_{i}, \mathbf{x}_{j}), (\mathbf{x}_{k}, \mathbf{x}_{l})).$$

Moreover, the preference kernel ensures that $f(x_i, x_j)$ is perfectly anti-correlated with the value of , $f(x_i, x_i)$, so that any sample from the GP will respect the anti-symmetric property of preference functions that $\tilde{f}(x_i, x_j) = -\tilde{f}(x_j, x_i)$. Indeed:

$$\operatorname{Corr}\left(f\left(x_{i}, x_{j}\right), f\left(x_{j}, x_{i}\right)\right) = \frac{k_{\operatorname{pref}}\left(\left(x_{i}, x_{j}\right), \left(x_{j}, x_{i}\right)\right)}{\sqrt{k_{\operatorname{pref}}\left(\left(x_{i}, x_{j}\right), \left(x_{i}, x_{j}\right)\right)}\sqrt{k_{\operatorname{pref}}\left(\left(x_{j}, x_{i}\right), \left(x_{j}, x_{i}\right)\right)}} = -1. \quad (F.1)$$

F.2 Conditioning on the value function offset

The mean function \hat{m}_{pref} and covariance function \hat{k}_{pref} of the GP prior on the preference function *f* can be computed from the mean function *m* and covariance function *k* of the GP on *g*

conditionally on $g(\mathbf{x}_0) = y_0$. Let \hat{k} be such that :

$$\hat{k}_{\text{pref}}((\boldsymbol{x}_i, \boldsymbol{x}_j), (\boldsymbol{x}_k, \boldsymbol{x}_l)) = \text{Cov}(f(\boldsymbol{x}_i, \boldsymbol{x}_j), f(\boldsymbol{x}_k, \boldsymbol{x}_l) | g(\boldsymbol{x}_0)).$$
(F.2)

The covariance can be decomposed:

$$\begin{aligned} \hat{k}_{\text{pref}}((\mathbf{x}_{i}, \mathbf{x}_{j}), (\mathbf{x}_{k}, \mathbf{x}_{l})) &= \text{Cov}(f(\mathbf{x}_{i}, \mathbf{x}_{j}), f(\mathbf{x}_{k}, \mathbf{x}_{l}) | g(\mathbf{x}_{0})) \\ &= \text{Cov}(g(\mathbf{x}_{i}) - g(\mathbf{x}_{j}), g(\mathbf{x}_{k}) - g(\mathbf{x}_{l}) | g(\mathbf{x}_{0})) \\ &= \text{Cov}(g(\mathbf{x}_{i}), g(\mathbf{x}_{k}) | g(\mathbf{x}_{0})) + \text{Cov}(g(\mathbf{x}_{j}), g(\mathbf{x}_{l}) | g(\mathbf{x}_{0})) - \\ &\text{Cov}(g(\mathbf{x}_{i}), g(\mathbf{x}_{l}) | g(\mathbf{x}_{0})) - \text{Cov}(g(\mathbf{x}_{j}), g(\mathbf{x}_{k}) | g(\mathbf{x}_{0})). \end{aligned}$$

From the formula for Gaussian conditionals, we have:

$$\operatorname{Cov}(g(\boldsymbol{x}_i), g(\boldsymbol{x}_k)|g(\boldsymbol{x}_0)) = k(\boldsymbol{x}_i, \boldsymbol{x}_j) - k(\boldsymbol{x}_i, \boldsymbol{x}_0)^\top k(\boldsymbol{x}_0, \boldsymbol{x}_0)^{-1} k(\boldsymbol{x}_0, \boldsymbol{x}_k).$$

By defining:

$$\hat{k}(\boldsymbol{x}_i, \boldsymbol{x}_j) = \operatorname{Cov}(g(\boldsymbol{x}_i), g(\boldsymbol{x}_k) | g(\boldsymbol{x}_0)),$$

we get:

$$\hat{k}_{\text{pref}}((x_i, x_j), (x_k, x_l)) = \hat{k}(x_i, x_j) + \hat{k}(x_k, x_l) - \hat{k}(x_i, x_l) - \hat{k}(x_k, x_j).$$
(F.3)

By definition, \hat{k}_{pref} generates valid covariance matrices. Similarly for the mean, by defining:

$$\hat{m}(\mathbf{x}_i) = \mathbb{E}[g(\mathbf{x}_i)|g(\mathbf{x}_0) = \mathbf{y}_0]$$

= $m(\mathbf{x}_i) + k(\mathbf{x}_i, \mathbf{x}_0)^\top k(\mathbf{x}_0, \mathbf{x}_0)^{-1} \mathbf{y}_0,$

we get:

$$\hat{m}_{\text{pref}}(\boldsymbol{x}_i, \boldsymbol{x}_j) = \mathbb{E}[f(\boldsymbol{x}_i, \boldsymbol{x}_j) | g(\boldsymbol{x}_0) = y_0]$$
$$= \mathbb{E}[g(\boldsymbol{x}_i) - g(\boldsymbol{x}_j) | g(\boldsymbol{x}_0) = y_0]$$
$$= \hat{m}(\boldsymbol{x}_i) - \hat{m}(\boldsymbol{x}_j).$$

G

G.1 Kernel approximation with random Fourier features

A shift-invariant kernel *k* can be approximated by means of the inner product of random features ϕ such that: $k(x, x') \sim \phi(x)^{\top} \phi(x')$. This method is known as the random Fourier features algorithm (or random kitchen sinks) (Rahimi and Recht, 2008), and was applied to GP by Lázaro-Gredilla et al. (2010) under the name Sparse Spectrum GP approximation. Briefly, the *m*-dimensional random feature is defined as $\phi(x) = \sqrt{2\alpha/m} \cos(Wx + \mathbf{b})$, where the $[W]_i$ are sampled from the the probability density corresponding to the normalized spectral density of the kernel, *b* is sampled uniformly in [0, 1] and α is a constant that depends on the kernel.

We have $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{p(\phi)} \left[\phi(\mathbf{x})^{\top} \phi(\mathbf{x}') \right]$, so that as the number of features *m* increases, the approximation gets more accurate. However, as the dimension of the input space increases, the number of spectral points required for a good approximation explodes.

Proof: The proof of this approximation can be found for example in Hernández-Lobato et al. (2014), and is reproduced here for completeness.

The kernel approximation is based on the Bochner's theorem.

Theorem 2 (*Bochner's theorem*) A continuous, shift-invariant kernel is positive definite if and only if it is the Fourier transform of a non-negative, finite measure.

As a result given some shift-invariant kernel k(x, x') = k(x - x', 0) there must exist an associated density s(w), known as its spectral density, which is the Fourier dual of k. This can be written as:

$$k(\mathbf{x}, \mathbf{x}') = \int e^{-i\mathbf{w}^{\top}(\mathbf{x} - \mathbf{x}')} s(\mathbf{w}) d\mathbf{w},$$

$$s(\mathbf{w}) = \frac{1}{(2\pi)^d} \int e^{i\mathbf{w}^{\top} \mathbf{\tau}} k(\mathbf{\tau}, \mathbf{0}) d\mathbf{\tau}.$$
(G.1)

Further, we can treat this measure as a probability density $p(w) = s(w)/\alpha$, where $\alpha = \int s(w)dw$ is the normalizing constant. Consequently, the kernel can be written as:

$$k(\mathbf{x},\mathbf{x}') = \alpha \mathbb{E}_{w} \left[e^{-i\mathbf{w}^{\top}(\mathbf{x}-\mathbf{x}')} \right].$$
(G.2)

And due to the symmetry of p(w) we can write the expectation as:

$$k(\mathbf{x}, \mathbf{x}') = \alpha \mathbb{E}_{w} \left[\frac{1}{2} \left(e^{-i\mathbf{w}^{\top}(\mathbf{x} - \mathbf{x}')} + e^{i\mathbf{w}^{\top}(\mathbf{x} - \mathbf{x}')} \right) \right]$$

= $\alpha \mathbb{E}_{w} \left[\cos \left(\mathbf{w}^{\top}\mathbf{x} - \mathbf{w}^{\top}\mathbf{x}' \right) \right].$ (G.3)

157

We can then note that $\int_0^{2\pi} \cos(a+2b)db = 0$ for any constant offset $a \in \mathbb{R}$. As a result, for *b* uniformly distributed between 0 and 2π we can write, by applying the sum of angles formula:

$$k(\mathbf{x},\mathbf{x}') = 2\alpha \mathbb{E}_{(w,b)} [\cos\left(\mathbf{w}^{\top}\mathbf{x} - b\right)\cos(\mathbf{w}^{\top}\mathbf{x}' + b)].$$
(G.4)

Finally, we can average over *m* weights and phases

$$k(\mathbf{x}, \mathbf{x}') = \frac{2\alpha}{m} \mathbb{E}_{(\mathbf{W}, \mathbf{b})} \left[\cos(\mathbf{W}\mathbf{x} + \mathbf{b})^{\top} \cos(\mathbf{W}\mathbf{x}' + \mathbf{b}) \right], \qquad (G.5)$$

where $[W]_i \sim p(w)$ and $[\mathbf{b}]_i \sim p(b)$ are stacked versions of the original random variables. The resulting quantity has the same expectation but results in a lower variance estimator. If we let $\phi(\mathbf{x}) = \sqrt{2\alpha/m} \cos(w\mathbf{x} + \mathbf{b})$ denote a random *m*-dimensional feature generated by this model we can also write the kernel as $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\phi} [\phi(\mathbf{x})^{\top} \phi(\mathbf{x}')]$.

G.2 Hilbert space kernel approximation

Recently, another method was proposed by Solin and Särkkä (2020), which aims at making the approximation as good as possible for a given rank (see Riutort-Mayol et al. (2020) for details about the practical implementation).

In this method, the kernel is approximated using a series expansion in terms of eigenfunctions of the Laplace operator on a rectangular domain $\Omega = [-L_1, L_1] \times \cdots \times [-L_d, L_d]$ (the search space are usually rectangular in Bayesian optimization). Briefly,

$$\widetilde{k}_{m}\left(\boldsymbol{x},\boldsymbol{x}'\right) = \sum_{j_{1},\dots,j_{d}=1}^{\widehat{m}} S\left(\sqrt{\lambda_{j_{1},\dots,j_{d}}}\right) \phi_{j_{1},\dots,j_{d}}\left(\boldsymbol{x}\right) \phi_{j_{1},\dots,j_{d}}\left(\boldsymbol{x}'\right),$$
(G.6)

where S is the spectral density of the kernel, $m = \hat{m}^d$ and with the eigenfunctions and eigenvalues:

$$\phi_{j_1,\dots,j_d}(\mathbf{x}) = \prod_{k=1}^d \frac{1}{\sqrt{L_k}} \sin\left(\frac{\pi j_k \left(\mathbf{x}_k + L_k\right)}{2L_k}\right),$$
(G.7)

and

$$\lambda_{j_1,\dots,j_d} = \sum_{k=1}^d \left(\frac{\pi j_k}{2L_k}\right)^2. \tag{G.8}$$

G.3 Equivalence between Bayesian linear model and finitedimensional approximation to Gaussian processes

Consider a Gaussian process prior $\mathcal{GP}(\mu(\cdot), k(\cdot, \cdot))$.

Assuming a finite approximation to the kernel $k(\mathbf{x}, \mathbf{x}') \sim \phi(\mathbf{x})^\top \phi(\mathbf{x}')$, the features $\phi(\mathbf{x})$ can be used to approximate the function f with a Bayesian linear model $f(\mathbf{x}) \sim \phi(\mathbf{x})^\top \boldsymbol{\omega}$, where $\boldsymbol{\omega} \sim \mathcal{N}(\mathbf{m}, \mathbf{V})$.

Proof: The proof comes from Hernández-Lobato et al. (2014).

Consider a Gaussian linear model $f(\mathbf{x}) = \phi(\mathbf{x})^\top \boldsymbol{\omega}$ where $\boldsymbol{\omega} \sim \mathcal{N}(0, \mathbf{I})$.

We assume that observations are sampled from a multivariate normal distribution with i.i.d. noise: $y|f(x) \sim \mathcal{N}(f(x), \sigma^2 I)$.

The posterior on ω is $\mathcal{N}(A\Phi^{\top}y, \sigma^2 A)$, where $A = (\Phi^{\top}\Phi + \sigma^2 I)^{-1}$, and $\Phi_{i,i} = \phi_i(x_i)$.

G.3. Equivalence between Bayesian linear model and finite-dimensional approximation to Gaussian processes

The predictive distribution over *g* is a Gaussian with mean and variance given by:

$$\mu(\mathbf{x}) = \phi(\mathbf{x})^{\top} (\Phi^{\top} \Phi + \sigma^2 \mathbf{I})^{-1} \Phi^{\top} \mathbf{y},$$

$$v(\mathbf{x}) = \sigma^2 \phi(\mathbf{x})^{\top} (\Phi^{\top} \Phi + \sigma^2 \mathbf{I})^{-1} \phi(\mathbf{x}).$$
(G.9)

By using the fact that $(I + PQ)^{-1}P = P(I + QP)^{-1}$, we get:

$$\mu(\mathbf{x}) = \phi(\mathbf{x})^{\top} \Phi^{\top} (\Phi \Phi^{\top} + \sigma^2 \mathbf{I})^{-1} \mathbf{y},$$
 (G.10)

and by applying the matrix inversion lemma:

$$v(\mathbf{x}) = \phi(\mathbf{x})^{\top} \phi(\mathbf{x}) - \phi(\mathbf{x})^{\top} \Phi^{\top} (\Phi \Phi^{\top} + \sigma^2 \mathbf{I})^{-1} \Phi \phi(\mathbf{x}).$$
(G.11)

By replacing the inner products with the corresponding kernel, we recover the predictive distribution of a Gaussian process (see appendice B).

Equivalence between the soft Copeland score distribution and the maximum latent value distribution

In Duel Thompson Sampling as originally described by (Gonzalez et al., 2017), a sample \tilde{f} is drawn from the posterior Gaussian process. The corresponding Condorcet winner is then computed. Here, we show that if we note \tilde{g} the corresponding sample from the posterior on the value function, then the corresponding Condorcet winner is the maximum of \tilde{g} . To show this, we show a stronger result:

$$\forall (\mathbf{x}, \mathbf{x}') \in \mathbf{x}^2, C(\mathbf{x}) = C(\mathbf{x}') \Leftrightarrow \tilde{g}(\mathbf{x}) = \tilde{g}(\mathbf{x}'). \tag{H.1}$$

Proof:

Consider a sample \tilde{g} from $p(g|\mathcal{D}_t)$. The corresponding sample from $p(f|\mathcal{D}_t)$ is \tilde{f} , such that $\tilde{f}(\boldsymbol{x}, \boldsymbol{x}') = \tilde{g}(\boldsymbol{x}) - \tilde{g}(\boldsymbol{x}')$

If x and x' are such that $\tilde{g}(x) = \tilde{g}(x')$, then $\forall y \in \mathcal{X}, \tilde{f}(x, y) = \tilde{f}(x', y)$, and since Φ is a bijective function on \mathbb{R} , we have that $\forall y \in \mathcal{X}, \Phi(\tilde{f}(x, y)) = \Phi(\tilde{f}(x', y))$, and hence C(x) = C(x').

Conversely, if $\exists (x, x') \in \mathcal{X}^2$ such that $\tilde{g}(x) \neq \tilde{g}(x')$, we can assume without loss of generality that $\tilde{g}(x) < \tilde{g}(x')$, and as a consequence $\forall y \in \mathcal{X}, \tilde{f}(x, y) < \tilde{f}(x', y)$, and since Φ is a strictly growing function on $\mathbb{R}, \forall y \in \mathcal{X}, \Phi(\tilde{f}(x, y)) < \Phi(\tilde{f}(x', y))$ and thus C(x) < C(x'). So by contrapositive, we get H.1.

The Laplace approximation for Gaussian process classification

The following description of Laplace approximation for Gaussian process classification is given for completeness. It is required in order to understand the computation of the gradient of the Binary Knowledge Gradient. We follow the reasoning presented in Rasmussen and Williams (2006) as well as Bishop (2006), to which we refer the reader for further details.

In Gaussian process classification, we assume that observations *c* at points $X = [x_1, \dots, x_t]$ are Bernoulli random variables with parameters $\mu_c(X) = \pi(f(X))$, where *f* is a latent function, and and π is an inverse link function. We choose the convention where *c* is either 0 or 1.

The predictive distribution at point *x* is given by:

$$p(c=1|\mathbf{x},\mathcal{D}_t) = \int p(c=1|f(\mathbf{x}))p(f(\mathbf{x})|\mathcal{D}_t)df(\mathbf{x}).$$
(I.1)

Since this integral is analytically intractable, we approximate $p(f(\mathbf{x})|\mathcal{D}_t)$ with a Gaussian distribution. Indeed, for a a random normally distributed random variable *z*:

$$\int \Phi(z) \mathcal{N}(z|\mu, \sigma^2) dz = \Phi\left(\frac{\mu}{\sqrt{1+\sigma^2}}\right).$$
(I.2)

A convenient way to do so is to note that:

$$p(f(\mathbf{x})|\mathscr{D}_t) = \int p(f(\mathbf{x})|f)p(f|\mathscr{D}_t)df,$$
(I.3)

where f is the vector of latent values at training points X. From the formula for Gaussian process posteriors, we have:

$$p(f(\mathbf{x})|\mathbf{f}) = \mathcal{N}\left(f(\mathbf{x})|\mathbf{k}^{\top}\mathbf{K}^{-1}\mathbf{f}, k(\mathbf{x}, \mathbf{x}) - \mathbf{k}^{\top}\mathbf{K}^{-1}\mathbf{k}\right),$$
(I.4)

where K = k(X, X) and k = k(X, x).

The second term in the integral $p(f|\mathscr{D}_t)$ is the posterior distribution of the latent value function at training points. By approximating it with a Gaussian distribution, then we could compute the integral in I.3, which would give us a Gaussian approximation for $p(f(\mathbf{x})|\mathscr{D}_t)$. This approximation would in turn allow us to compute I.1.

I.1 Principle of the Laplace approximation

We start by finding a Gaussian approximation of $p(f|\mathcal{D}_t)$. In general, for a random variable z whose probability density function is $p(z) = \frac{f(z)}{Z}$, we can use a second a second order Taylor expansion around the mode z_0 of the distribution (where the gradient vanishes) so that:

$$\ln f(z) \sim \ln f(z_0) - \frac{1}{2} (z - z_0)^\top H(z - z_0),$$
 (I.5)

161

where *H* is the negative of the Hessian of f at z_0 .

Taking the exponential and computing the appropriate normalization coefficients, we have:

$$p(z) \sim \frac{|\mathbf{H}|^{\frac{1}{2}}}{(2\pi)^{\frac{D}{2}}} \exp\left(-\frac{1}{2}(z-z_0)^{\top} \mathbf{H}(z-z_0)\right) = \mathcal{N}(z|z_0, \mathbf{H}^{-1})$$
(I.6)

This method of approximating a probability density function with a Gaussian probability density function is called the Laplace approximation.

I.2 Gaussian approximation of the posterior

In order to find a Gaussian approximation of $p(f|\mathscr{D}_t)$, we thus need to compute its mode and its Hessian. By using Bayes' rule, we have:

$$\ln p(f|\mathcal{D}_t) = \ln p(f) + \ln p(\mathcal{D}_t|f). \tag{I.7}$$

The prior term is:

$$\ln p(f) = -\frac{1}{2} f^{\top} K^{-1} f - \frac{t}{2} \ln(2\pi) - \frac{1}{2} \ln|K|.$$
 (I.8)

The likelihood term is:

$$\ln p(\mathscr{D}_t | f) = \ln \left(\prod_{i=1}^t \pi(f_i)^{c_i} (1 - \pi(f_i))^{(1 - c_i)} \right)$$

= $\sum_{i=1}^t \ln \left(\pi(f_i)^{c_i} (1 - \pi(f_i))^{(1 - c_i)} \right).$ (I.9)

So the gradient of the the log-posterior is:

$$\nabla_f \ln p(f|\mathscr{D}_t) = \nabla_f \ln p(\mathscr{D}_t|f) - K^{-1}f, \qquad (I.10)$$

whereas the Hessian is:

$$\nabla_f^2 \ln p(f|\mathscr{D}_t) = \nabla_f^2 \ln p(\mathscr{D}_t|f) - K^{-1}, \qquad (I.11)$$

where ∇^2 refers to the Hessian matrix.

We introduce $W = -\nabla_f^2 \ln p(\mathcal{D}_t | f)$, which is a diagonal matrix since conditionally on f, observations are independent.

$$\frac{\partial \ln p(c_i|f_i)}{\partial f_i} = \frac{\pi'(f_i)(c_i - \pi(f_i))}{\pi(f_i)(1 - \pi(f_i))}.$$
(I.12)

In the case where the link function is the cumulative normal distribution:

$$\frac{\partial \ln p(c_i|f_i)}{\partial f_i} = \frac{(2c_i - 1)\phi(f_i)}{\Phi((2c_i - 1)f_i)},$$
(I.13)

and

$$\frac{\partial^2 \ln p(c_i|f_i)}{\partial f_i^2} = -\frac{\phi(f_i)^2}{\Phi((2c_i - 1)f_i)^2} - \frac{(2c_i - 1)f_i\phi(f_i)}{\Phi((2c_i - 1)f_i)}.$$
(I.14)

The mode f_0 satisfies the condition $\nabla_f \ln p(f|\mathscr{D}_t) = 0$, so $f_0 = K \nabla_f \ln p(\mathscr{D}_t|f)$. The mode is usually found using the Newton-Raphson method.

The approximate posterior distribution is :

$$p(f|\mathcal{D}_t) = \mathcal{N}(f_0, (K^{-1} + W)^{-1}).$$
(I.15)

I.3 Approximate predictive distribution

By combining equation with equation I.15, we get:

$$\mathbb{E}_f[f(\mathbf{x})|\mathscr{D}_t] = \mathbf{k}^\top \mathbf{K}^{-1} f_0, \tag{I.16}$$

and:

$$\mathbb{V}_f[f(\mathbf{x})|\mathscr{D}_t] = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}^\top (\mathbf{K} + \mathbf{W}^{-1})^{-1} \mathbf{k}.$$
 (I.17)

Benchmarks

Name	D	Kernel	Space
Ackley	2	Matérn 3/2	$[-32.768, 32.768]^2$
Beale	2	SE-ARD	$[-4.5, 4.5]^2$
Bohachevsky	2	SE-ARD	$[-100, 100]^2$
Three-Hump Camel	2	Matérn 5/2	$[-5,5]^2$
Six-Hump Camel	2	SE-ARD	$[-3,3] \times [-2,2]$
Colville	4	Matérn 5/2	$[-10, 10]^4$
Cross-in-Tray	2	Matérn 5/2	$[-10, 10]^2$
Dixon-Price	2	Matérn 5/2	$[-5, 5]^2$
Drop-Wave	2	Matérn 3/2	$[-5.12, 5.12]^2$
Eggholder	2	SE-ARD	$[-512, 512]^2$
Forrester et al (2008)	1	SE-ARD	[0,1]
Goldstein-Price	2	SE-ARD	$[-2, 2]^2$
Griewank	2	SE-ARD	$[-600, 600]^2$
Gramacy and Lee (2012)	1	SE-ARD	[0.5, 2.5]
Hartmann 3-D	3	SE-ARD	$[0, 1]^3$
Hartmann 4D	4	SE-ARD	$[0, 1]^4$
Hartmann 6D	6	SE-ARD	$[0, 1]^6$
Holder	2	SE-ARD	$[-10, 10]^2$
Langer	2	Matérn 3/2	$[0, 10]^2$
Levy	2	SE-ARD	$[-10, 10]^2$
Levy N.13	2	Matérn 5/2	$[-10, 10]^2$
Perm 0,d, β	2	SE-ARD	$[-2, 2]^2$
Perm d, β	2	SE-ARD	$[-2, 2]^2$
Powell	4	SE-ARD	$[-4, 5]^4$
Rosenbrock	2	SE-ARD	$[-2.048, 2.048]^2$
Rotated Hyper-Ellipsoid	2	Matérn 3/2	$[-65.536, 65.536]^2$
Schaffer n4	2	Matérn 3/2	$[-100, 100]^2$
Schwefel	2	SE-ARD	$[-500, 500]^2$
Shekel	4	SE-ARD	$[0, 10]^4$
Schubert	2	Matérn 3/2	$[0, 10]^2$
Sphere	2	SE-ARD	$[-5.12, 5.12]^2$

J

Name	D	Kernel	Space
Sum Squares	2	SE-ARD	$[-10, 10]^2$
Trid	2	SE-ARD	$[-4, 4]^2$
Ursem Waves	2	SE-ARD	$[-1.2, 1.2] \times [-0.9, 1.2]$

 Table J.1: Benchmark functions in Bayesian optimization experiments.
1.1		
	The human visual system	3
1.2	Signal processing in the retina	5
1.3	Phenotype of retinitis pigmentosa	7
1.4	Stages of AMD	9
1.5	Electronic visual prosthetics designs	14
1.6	Phosphene modulation with pulses amplitude	18
1.7	Phosphene modulation with pulses frequency	19
1.8	An example of image processing in the Argus II retinal implant	21
1.9	New generation retinal prostheses: PRIMA and POLYRETINA	27
1.10	Electrodes crosstalk	29
1.11	Distortion due to passing fibers activation	31
1.12	Signal distortion due to unspecific cell types stimulation in simulated prosthetic	
	vision	32
1.13	Simulated prosthetic vision in Argus II users	34
1.14	Recordings-based closed-loop optimization strategy	35
1.15	Retinal prosthetic encoder	37
1.16	Principle of Human-in-the-loop optimization	41
1.17	Retinal prosthetic encoder optimization	44
2.1	Encoder parameterization.	49
2.2	Principle of hyperparameters transfer.	55
2.3	Effect of the parameterization	58
2.3 2.4	Effect of the parameterization	58 59
 2.3 2.4 2.5 	Effect of the parameterization	58 59 61
 2.3 2.4 2.5 2.6 	Effect of the parameterizationOptimization sequenceRelationship between MSE and preference between encoders.Visual acuity measurements following preference-based optimization.	58 59 61 62
 2.3 2.4 2.5 2.6 2.7 	Effect of the parameterizationOptimization sequenceRelationship between MSE and preference between encoders.Visual acuity measurements following preference-based optimization.Performance-based optimization	58 59 61 62 63
 2.3 2.4 2.5 2.6 2.7 2.8 	Effect of the parameterizationOptimization sequenceRelationship between MSE and preference between encoders.Visual acuity measurements following preference-based optimization.Performance-based optimizationModel misspecification	58 59 61 62 63 64
 2.3 2.4 2.5 2.6 2.7 2.8 2.9 	Effect of the parameterizationOptimization sequenceRelationship between MSE and preference between encoders.Visual acuity measurements following preference-based optimization.Performance-based optimizationModel misspecificationOptimization with a misspecified model.	58 59 61 62 63 64 65
 2.3 2.4 2.5 2.6 2.7 2.8 2.9 2.10 	Effect of the parameterizationOptimization sequenceRelationship between MSE and preference between encoders.Visual acuity measurements following preference-based optimization.Performance-based optimizationModel misspecificationOptimization with a misspecified model.Optimization with a misspecified model.	 58 59 61 62 63 64 65 70
 2.3 2.4 2.5 2.6 2.7 2.8 2.9 2.10 2.11 	Effect of the parameterizationOptimization sequenceRelationship between MSE and preference between encoders.Visual acuity measurements following preference-based optimization.Performance-based optimizationModel misspecificationOptimization with a misspecified model.Optimization with a misspecified model.Remote vs lab.	 58 59 61 62 63 64 65 70 71
 2.3 2.4 2.5 2.6 2.7 2.8 2.9 2.10 2.11 3.1 	Effect of the parameterizationOptimization sequenceRelationship between MSE and preference between encoders.Visual acuity measurements following preference-based optimization.Performance-based optimizationModel misspecificationOptimization with a misspecified model.Optimization with a misspecified model.Remote vs lab.Example of Gaussian process regression	58 59 61 62 63 64 65 70 71 77
 2.3 2.4 2.5 2.6 2.7 2.8 2.9 2.10 2.11 3.1 3.2 	Effect of the parameterizationOptimization sequenceRelationship between MSE and preference between encoders.Visual acuity measurements following preference-based optimization.Performance-based optimizationModel misspecificationOptimization with a misspecified model.Optimization with a misspecified model.Remote vs lab.Example of Gaussian process regressionGaussian process regression	 58 59 61 62 63 64 65 70 71 77 78
 2.3 2.4 2.5 2.6 2.7 2.8 2.9 2.10 2.11 3.1 3.2 3.3 	Effect of the parameterization	 58 59 61 62 63 64 65 70 71 77 78 79

List of Figures

3.5	Conditional Preference kernel	31
3.6	Effect of value function conditioning on preference learning, as a function of	
	training set size	32
3.7	Kernels comparison	33
3.8	Effect of kernel hyperparameters on the performance of Bayesian optimization	34
3.9	Bayesian optimization with continuous outputs	35
3.10	Uncertainty decomposition.	38
3.11	Active learning.	39
3.12	Identification step) 0
3.13	Comparison of acquisition functions for binary BO.	92
3.14	Comparison between UCB in the latent and the probability space) 3
3.15	Acquisition rules behavior in BBO	9 5
3.16	Pathological behavior of the original Expected Improvement acquisition func-	
	tion by Brochu et al. (2008, 2010a)) 7
3.17	Maximally Uncertain Challenge)2
3.18	Behavior of acquisition functions of the 'challenge' class)3
3.19	BBO trajectories)4
3.20	PBO trajectories)6
3.21	Batch preferential Bayesian optimization)7
3.22	Approximate GP sampling	10
3.23	Kernel approximation	11
3.24	Sampling from a preference GP	13
3.25	Comparison of approximate sampling methods	14
3.26	Detailed comparison of acquisition functions in binary BO	16
3.27	Detailed comparison of acquisition functions in preferential BO	17
4.1	Contextual binary Knowledge Gradient	22
4.2	Binary Bayesian optimization with adaptive contexts	25
4.3	Preference learning with adaptive contexts	26
4.4	Visual acuity model	<u>29</u>
4.5	Refractive error correction	30
4.6	Detailed results of performance comparison between acquisition functions in	
	contextual Binary Bayesian optimization. Each entry (i,j) corresponds to the frac-	
	tion of benchmarks functions for which i beats j according to the Mann-Whitney	
	U test at $\alpha = 5 \times 10^{-4}$ significance based either on the best value found (A) or	
	the Area Under the Curve (B)	33
4.7	Detailed results of performance comparison between acquisition functions in	
	contextual Preferential Bayesian optimization. Each entry (i,j) corresponds to	
	the fraction of benchmarks functions for which i beats j according to the Mann-	
	Whitney U test at $\alpha = 5 \times 10^{-4}$ significance based either on the best value found	
	(A) or the Area Under the Curve (B)	34

2.1	Search space in the Human-in-the-loop optimization of retinal prosthetic encoder. 54
3.1	List of acquisition rules for binary Bayesian optimization
3.2	List of acquisition rules for Preferential Bayesian Optimization
3.3	Comparison of acquisition rules on benchmarks in BBO
3.4	Comparison of acquisition rules on benchmarks in PBO
3.5	Comparison of acquisition rules on benchmarks in batch PBO
4.1	Comparison of acquisition rules on benchmarks for cBBO
4.2	Comparison of acquisition rules on benchmarks for cPBO
4.3	Comparison of acquisition rules on the refractive error correction experiment 130
J.1	Benchmarks

Acronyms

- AMD Age-related Macular Degeneration.
- **AMPA** α -amino-3-hydroxy-5-methyl-4-isoxazole propionic acid.
- **BBO** Binary Bayesian optimization.
- EI Expected improvement.
- EIIG Expected Improvement with Information Gain.
- GCL Ganglion Cell Layer.
- INL Inner Nuclear Layer.
- **IPL** Inner Plexiform Layer.
- KG Knowledge gradient.
- LGN Lateral Geniculate Nucleus.
- MUC Maximally Uncertain Challenge.
- NVAMD Neovascular Age-related Macular Degeneration.
- **OPL** Outer Plexiform Layer.
- **PBO** Preferential Bayesian optimization.
- **RGC** Retina Ganglion Cells.
- **RP** Retinitis Pigmentosa.
- **TS** Thompson Sampling.
- UCB Upper Credible Bound.
- V1 Primary Visual Cortex.

- M. Abramian, N. H. Lovell, J. W. Morley, G. J. Suaning, and S. Dokos. Activation of retinal ganglion cells following epiretinal electrical stimulation with hexagonally arranged bipolar electrodes. *J. Neural Eng.*, 8(3):35004, 2011. [Cited on page 28.]
- W. F. Agnew, D. B. McCreery, T. G. Yuen, and L. A. Bullara. MK-801 protects against neuronal injury induced by electrical stimulation. *Neuroscience*, 52(1):45–53, 1993. [Cited on page 25.]
- A. K. Ahuja, J. D. Dorn, A. Caspi, M. J. McMahon, G. Dagnelie, L. DaCruz, P. Stanga, M. S. Humayun, and R. J. Greenberg. Blind subjects implanted with the Argus II retinal prosthesis are able to improve performance in a spatial-motor task. *Br. J. Ophthalmol.*, 95(4):539–543, 2011. [Cited on pages 23 and 48.]
- M. J. I. Airaghi Leccardi. *Development of a new visual prosthesis for preclinical studies on artificial vision*. PhD thesis, EPFL, 2020. [Cited on pages 3 and 14.]
- N. Aït-Ali, R. Fridlich, G. Millet-Puel, E. Clérin, F. Delalande, C. Jaillard, F. Blond, L. Perrocheau, S. Reichman, L. C. Byrne, A. Olivier-Bandini, J. Bellalou, E. Moyse, F. Bouillaud, X. Nicol, D. Dalkara, A. Van Dorsselaer, J. A. Sahel, and T. Léveillard. Rod-derived cone viability factor promotes cone survival by stimulating aerobic glycolysis. *Cell*, 161(4):817–832, 2015. [Cited on page 7.]
- A. I. Al-Kalbani, M. R. Yuce, and J. M. Redoute. Safe SAR levels in inductively powered brain implanted visual prostheses. *IEEE Int. Symp. Electromagn. Compat.*, 2012. [Cited on page 17.]
- V. Anand, B. Duffy, Z. Yang, N. S. Dejneka, A. M. Maguire, and J. Bennett. A deviant immune response to viral proteins and transgene product is generated on subretinal administration of adenovirus and adeno-associated virus. *Mol. Ther.*, 5(2):125–132, 2002. [Cited on page 12.]
- M. R. Antognazza, M. Di Paolo, D. Ghezzi, M. Mete, S. Di Marco, J. F. Maya-Vetencourt, R. Maccarone, A. Desii,
 F. Di Fonzo, M. Bramini, A. Russo, L. Laudato, I. Donelli, M. Cilli, G. Freddi, G. Pertile, G. Lanzani, S. Bisti, and
 F. Benfenati. Characterization of a Polymer-Based, Fully Organic Prosthesis for Implantation into the Subretinal
 Space of the Rat. *Adv. Healthc. Mater.*, 5(17):2271–2282, 2016. [Cited on page 16.]
- R. Antonova, A. Rai, and C. G. Atkeson. Deep Kernels for Optimizing Locomotion Controllers. In *ICML*, 2017. [Cited on page 143.]
- L. N. Ayton, P. J. Blamey, R. H. Guymer, C. D. Luu, D. A. Nayagam, N. C. Sinclair, M. N. Shivdasani, J. Yeoh, M. F. McCombe, R. J. Briggs, N. L. Opie, J. Villalobos, P. N. Dimitrov, M. Varsamidis, M. A. Petoe, C. D. McCarthy, J. G. Walker, N. Barnes, A. N. Burkitt, C. E. Williams, R. K. Shepherd, and P. J. Allen. First-in-human trial of a novel suprachoroidal retinal prosthesis. *PLoS One*, 9(12):1–26, 2014. [Cited on pages 14, 24, 24, and 56.]
- L. N. Ayton, J. F. Rizzo, I. L. Bailey, A. Colenbrander, G. Dagnelie, D. R. Geruschat, P. C. Hessburg, C. D. McCarthy, M. A. Petoe, G. S. Rubin, P. R. Troyk, and H. International. Harmonization of Outcomes and Vision Endpoints in Vision Restoration Trials: Recommendations from the International HOVER Taskforce. *Transl. Vis. Sci. Technol.*, 9 (8):25, 2020. [Cited on pages 21 and 22.]
- J. Azimi, A. Jalali, and X. Z. Fern. Hybrid batch bayesian optimization. *Proc. 29th Int. Conf. Mach. Learn. ICML 2012*, 2(1):1215–1222, 2012. [Cited on page 108.]

- P. Bach-y Rita and C. C. Collins. Sensory Substitution Systems Using the Skin for the Input to the Brain. J. Audio Eng. Soc, 19(5):427–429, 1971. [Cited on page 16.]
- P. Bach-y Rita and S. W. Kercel. Sensory substitution and the human-machine interface. *Trends Cogn. Sci.*, 7(12): 541–546, 2003. [Cited on page 16.]
- P. Bach-y rita, C. C. Collins, F. A. Saunders, B. White, and L. Scadden. Vision Substitution by Tactile Image Projection. *Nature*, 221(5184):963–964, 1969. [Cited on page 16.]
- P. Bach-y Rita, K. A. Kaczmarek, M. E. Tyler, and J. Garcia-Lara. Form perception with a 49-point electrotactile stimulus array on the tongue: A technical note. *J. Rehabil. Res. Dev.*, 35(4):427–430, 1998. [Cited on page 16.]
- T. Baden, P. Berens, K. Franke, M. Román Rosón, M. Bethge, and T. Euler. The functional diversity of retinal ganglion cells in the mouse. *Nature*, 529(7586):345–350, 2016. [Cited on page 4.]
- T. Baden, T. Euler, and P. Berens. Understanding the retinal basis of vision across species. *Nat. Rev. Neurosci.*, pages 1–16, 2019. [Cited on page 2.]
- M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy. BoTorch: Programmable Bayesian Optimization in PyTorch. *arXiv*, 2019a. [Cited on page 113.]
- M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. *Adv. Neural Inf. Process. Syst.*, 33, 2019b. [Cited on page 113.]
- A. C. Barber, C. Hippert, Y. Duran, E. L. West, J. W. Bainbridge, K. Warre-Cornish, U. F. Luhmann, J. Lakowski, J. C. Sowden, R. R. Ali, and R. A. Pearson. Repair of the degenerate retina by photoreceptor transplantation. *Proc. Natl. Acad. Sci. U. S. A.*, 110(1):354–359, 2013. [Cited on page 12.]
- L. Bareket, A. Barriga-Rivera, M. P. Zapf, N. H. Lovell, and G. J. Suaning. Progress in artificial vision through suprachoroidal retinal implants. *J. Neural Eng.*, 14(4):45002, 2017. [Cited on page 17.]
- N. Barnes, A. F. Scott, P. Lieby, M. A. Petoe, C. McCarthy, A. Stacey, L. N. Ayton, N. C. Sinclair, M. N. Shivdasani, N. H. Lovell, H. J. McDermott, and J. G. Walker. Vision function testing for a suprachoroidal retinal prosthesis: Effects of image filtering. *J. Neural Eng.*, 13(3):1–15, 2016. [Cited on pages 20, 20, 47, and 51.]
- J. R. Bartlett, R. W. Doty, sr., B. B. Lee, N. Negrão, and W. H. Overman, jr. Deleterious Effects of Prolonged Electrical Excitation of Striate Cortex in Macaques. *Brain. Behav. Evol.*, 14(1-2):46–66, 1977. [Cited on page 25.]
- H. A. Baseler, A. Gouws, K. V. Haak, C. Racey, M. D. Crossland, A. Tufail, G. S. Rubin, F. W. Cornelissen, and A. B. Morland. Large-scale remapping of visual cortex is absent in adult humans with macular degeneration. *Nat. Neurosci.*, 14(5):649–657, 2011. [Cited on pages 10 and 10.]
- P. J. Basser and B. J. Roth. New Currents in Electrical Stimulation of Excitable Tissues. *Annu. Rev. Biomed. Eng.*, 2(1): 377–397, 2000. [Cited on page 13.]
- K. Basu and S. Ghosh. Adaptive Rate of Convergence of Thompson Sampling for Gaussian Process Optimization. *arXiv*, pages 1–29, 2017. [Cited on pages 42 and 122.]
- M. Becker, R. Eckmiller, and R. Huenermann. Psychophysical test of a tunable retina encoder for retina implants. In *Int. Jt. Conf. Neural Networks (IJCNN'99), July 10, 1999 July 16, 1999*, volume 1, pages 192–195. IEEE, 1999. [Cited on page 66.]
- M. R. Behrend, A. K. Ahuja, M. S. Humayun, J. D. Weiland, and R. H. Chow. Selective labeling of retinal ganglion cells with calcium indicators by retrograde loading in vitro. *J. Neurosci. Methods*, 179(2):166–172, 2009. [Cited on page 29.]
- A. Bellet, A. Habrard, and M. Sebban. A Survey on Metric Learning for Feature Vectors and Structured Data. *arXiv*, 2013. [Cited on page 143.]

- A. Benavoli, D. Azzimonti, and D. Piga. Skew Gaussian Processes for Classification. *arXiv*, 2020a. [Cited on pages 77, 105, 105, and 105.]
- A. Benavoli, D. Azzimonti, and D. Piga. Preferential Bayesian optimisation with Skew Gaussian Processes. *arXiv*, pages 1–18, 2020b. [Cited on pages 77, 95, 101, 101, 101, and 108.]
- Y. Bengio, O. Delalleau, and N. Le Roux. The curse of highly variable functions for local kernel machines. Adv. Neural Inf. Process. Syst., pages 107–114, 2005. [Cited on page 141.]
- V. Bengs, R. Busa-Fekete, A. E. Mesaoudi-Paul, and E. Hüllermeier. Preference-based Online Learning with Dueling Bandits: A Survey. J. Mach. Learn. Res., 22:1–108, 2021. [Cited on pages 53, 94, and 115.]
- J. Bennett. Immune response following intraocular delivery of recombinant viral vectors. *Gene Ther.*, 10(11 SPEC.): 977–982, 2003. [Cited on page 12.]
- J. O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer Series in Statistics. Springer New York, New York, NY, 2nd edition, 1985. [Cited on page 53.]
- W. Berger, B. Kloeckener-Gruissem, and J. Neidhardt. The molecular basis of human retinal and vitreoretinal diseases. *Prog. Retin. Eye Res.*, 29(5):335–375, 2010. [Cited on page 6.]
- G. K. Bergey. Neurostimulation in the treatment of epilepsy. Exp. Neurol., 244:87–95, 2013. [Cited on page 13.]
- J. Bergstra and Y. Bengio. Random Search for Hyper-Parameter Optimization. J. Mach. Learn. Res., pages 281–305, 2012. [Cited on pages 40 and 100.]
- F. Berkenkamp, A. P. Schoellig, and A. Krause. No-regret bayesian optimization with unknown hyperparameters. J. Mach. Learn. Res., 20:1–24, 2019. [Cited on page 136.]
- A. Berntson and W. R. Taylor. Response characteristics and receptive field widths of on-bipolar cells in the mouse retina. *J. Physiol.*, 524(3):879–889, 2000. [Cited on page 3.]
- E. L. Berson. Retinitis pigmentosa: The Friedenwald lecture. *Investig. Ophthalmol. Vis. Sci.*, 34(5):1659–1676, 1993. [Cited on page 6.]
- M. Bertelsen, H. Jensen, J. F. Bregnhoj, and T. Rosenberg. Prevalence of generalized retinal dystrophy in Denmark. *Ophthalmic Epidemiol.*, 21(4):217–223, 2014. [Cited on page 6.]
- M. Beyeler. Biophysical model of axonal stimulation in epiretinal visual prostheses. *arXiv*, 2018. [Cited on pages 30, 30, and 46.]
- M. Beyeler, G. Boynton, I. Fine, and A. Rokem. pulse2percept: A Python-based simulation framework for bionic vision. In *Proc. 16th Python Sci. Conf.*, pages 81–88. SciPy, 2017a. [Cited on pages 30, 33, 33, 34, 48, 50, 57, 65, 135, and 143.]
- M. Beyeler, A. Rokem, G. M. Boynton, and I. Fine. Learning to see again: biological constraints on cortical plasticity and the implications for sight restoration technologies. *J. Neural Eng.*, 14(5):051003, 2017b. [Cited on pages 10, 11, 21, 29, and 69.]
- M. Beyeler, G. M. Boynton, I. Fine, and A. Rokem. Model-Based Recommendations for Optimal Surgical Placement of Epiretinal Implants. In *MICCAI*, 2019a. [Cited on pages 29, 30, 30, and 31.]
- M. Beyeler, D. Nanduri, J. D. Weiland, A. Rokem, G. M. Boynton, I. Fine, G. M. Boynton, and I. Fine. A model of ganglion axon pathways accounts for percepts elicited by retinal implants. *bioRxiv*, 9(1):1–6, 2019b. [Cited on pages 28, 33, 33, 34, 46, 47, 47, 48, 48, 51, 51, 54, 64, 135, 135, 139, and 141.]
- C. M. Bishop. Pattern Recognition and Machine Learning. Springer, 2006. [Cited on pages 150 and 161.]
- R. Blendowske. Unaided Visual Acuity and Blur. *Optom. Vis. Sci.*, 92(6):e121–e125, 2015. [Cited on pages 128 and 128.]

- E. Bloch and L. da Cruz. The Argus II Retinal Prosthesis System. In *Prosthesis*. IntechOpen, 2019. [Cited on pages 14, 21, and 57.]
- D. Boinagrov, S. Pangratz-Fuehrer, G. Goetz, and D. Palanker. Selectivity of direct and network-mediated stimulation of the retinal ganglion cells with epi-, sub- and intraretinal electrodes. *J. Neural Eng.*, 11(2), 2014. [Cited on page 30.]
- M. Boulton, P. Moriarty, J. Jarvis-Evans, and B. Marcyniuk. Regional variation and age-related changes of lysosomal enzymes in the human retinal pigment epithelium. *Br. J. Ophthalmol.*, 78(2):125–129, 1994. [Cited on page 5.]
- R. R. Bourne, G. A. Stevens, R. A. White, J. L. Smith, S. R. Flaxman, H. Price, J. B. Jonas, J. Keeffe, J. Leasher, K. Naidoo, K. Pesudovs, S. Resnikoff, and H. R. Taylor. Causes of vision loss worldwide, 1990-2010: A systematic analysis. *Lancet Glob. Heal.*, 1(6):339–349, 2013. [Cited on page 6.]
- R. R. Bourne, J. B. Jonas, A. M. Bron, M. V. Cicinelli, A. Das, S. R. Flaxman, D. S. Friedman, J. E. Keeffe, J. H. Kempen, J. Leasher, H. Limburg, K. Naidoo, K. Pesudovs, T. Peto, J. Saadine, A. J. Silvester, N. Tahhan, H. R. Taylor, R. Varma, T. Y. Wong, and S. Resnikoff. Prevalence and causes of vision loss in high-income countries and in Eastern and Central Europe in 2015: Magnitude, temporal trends and projections. *Br. J. Ophthalmol.*, 102 (5):575–585, 2018. [Cited on page 6.]
- D. H. Brainard. The Psychophysics Toolbox. Spat. Vis., 10(4):433-436, 1997. [Cited on page 69.]
- M. E. Brelén, F. Duret, B. Gérard, J. Delbeke, and C. Veraart. Creating a meaningful visual perception in blind volunteers by optic nerve stimulation. *J. Neural Eng.*, 2(1):S22–S28, 2005. [Cited on page 14.]
- M. E. Brelén, P. De Potter, M. Gersdorff, G. Cosnard, C. Veraart, and J. Delbeke. Intraorbital implantation of a stimulating electrode for an optic nerve visual prosthesis. *J. Neurosurg.*, 104(4):593–597, 2006. [Cited on page 14.]
- M. E. Brelén, V. Vince, B. Gérard, C. Veraart, and J. Delbeke. Measurement of Evoked Potentials after Electrical Stimulation of the Human Optic Nerve. *Investig. Opthalmology Vis. Sci.*, 51(10):5351, 2010. [Cited on page 14.]
- T. Breuninger, C. Puller, S. Haverkamp, and T. Euler. Chromatic bipolar cell pathways in the mouse retina. *J. Neurosci.*, 31(17):6504–6517, 2011. [Cited on page 3.]
- G. S. Brindley and W. S. Lewin. The sensations produced by electrical stimulation of the visual cortex. *J. Physiol.*, 196(2):479–93, 1968. [Cited on pages 15 and 28.]
- E. Brochu, N. D. Freitas, and A. Ghosh. Active preference learning with discrete choice data. *Adv. Neural Inf. Process. Syst.*, pages 409–416, 2008. [Cited on pages 96, 96, 97, and 167.]
- E. Brochu, T. Brochu, and N. Freitas. A Bayesian interactive optimization approach to procedural animation design. *Comput. Animat.* 2010 - ACM SIGGRAPH / Eurographics Symp. Proceedings, SCA 2010, pages 103–112, 2010a. [Cited on pages 43, 52, 75, 96, 97, 100, and 167.]
- E. Brochu, V. M. Cora, and N. de Freitas. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. *ArXiv*, 2010b. [Cited on pages 51, 51, 52, 72, 74, 75, 96, 96, 101, 103, 105, and 123.]
- S. B. Brummer, L. S. Robblee, and F. T. Hambrecht. Criteria for Selecting Electrodes for Electrical Stimulation: Theoretical and Practical Considerations. *Ann. N. Y. Acad. Sci.*, 405(1):159–171, 1983. [Cited on page 25.]
- S. Bubeck, R. Munos, and G. Stoltz. Pure exploration in multi-armed bandits problems. *Lect. Notes Comput. Sci.* (*including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics*), 5809 LNAI:23–37, 2009. [Cited on page 115.]
- H. Buch, T. Vinding, M. La Cour, M. Appleyard, G. B. Jensen, and N. V. Nielsen. Prevalence and Causes of Visual Impairment and Blindness among 9980 Scandinavian Adults: The Copenhagen City Eye Study. *Ophthalmology*, 111(1):53–61, 2004. [Cited on page 6.]
- C. Büchel, C. Price, R. S. Frackowiak, and K. Friston. Different activation patterns in the visual cortex of late and congenitally blind subjects. *Brain*, 121(3):409–419, 1998. [Cited on page 10.]

- L. A. Bullara, D. B. McCreery, W. F. Agnew, and T. G. Yuen. Comparison of neural damage induced by electrical stimulation with faradaic and capacitor electrodes. *Ann. Biomed. Eng.*, 16(5):463–481, 1988. [Cited on page 25.]
- V. Busskamp, S. Picaud, J. A. Sahel, and B. Roska. Optogenetic therapy for retinitis pigmentosa. *Gene Ther.*, 19(2): 169–175, 2012. [Cited on page 12.]
- C. Cai, Q. Ren, N. J. Desai, J. F. Rizzo, and S. I. Fried. Response variability to high rates of electric stimulation in retinal ganglion cells. *J. Neurophysiol.*, 106(1):153–162, 2011. [Cited on page 31.]
- C. Cai, P. Twyford, and S. Fried. The response of retinal neurons to high-frequency stimulation. *J. Neural Eng.*, 10 (3):036009, 2013. [Cited on page 30.]
- D. Calandriello, L. Carratino, A. Lazaric, M. Valko, and L. Rosasco. Gaussian Process Optimization with Adaptive Sketching: Scalable and No Regret. In *Proc. Mach. Learn. Res.*, volume 99, pages 1–25, 2019. [Cited on page 109.]
- P. A. Campochiaro and T. A. Mir. The mechanism of cone cell death in Retinitis Pigmentosa. *Prog. Retin. Eye Res.*, 62:24–37, 2018. [Cited on page 6.]
- C. Capelle, C. Trullemans, P. Arno, and C. Veraart. A real-time experimental prototype for enhancement of vision rehabilitation using auditory substitution. *IEEE Trans. Biomed. Eng.*, 45(10):1279–1293, 1998. [Cited on page 16.]
- A. Carkeet, L. Lee, J. R. Kerr, and M. M. Keung. The slope of the psychometric function for Bailey-Lovie letter charts: Defocus effects and impilications for modeling letter-by-letter scores. *Optom. Vis. Sci.*, 78(2):113–121, 2001. [Cited on page 128.]
- M. L. Carlson, C. L. Driscoll, R. H. Gifford, and S. O. McMenomey. Cochlear implantation: Current and future device options. *Otolaryngol. Clin. North Am.*, 45(1):221–248, 2012. [Cited on page 13.]
- E. Castaldi, G. M. Cicchini, L. Cinelli, L. Biagi, S. Rizzo, and M. C. Morrone. Visual BOLD Response in Late Blind Subjects with Argus II Retinal Prosthesis. *PLoS Biol.*, 14(10):1–19, 2016. [Cited on pages 21, 21, and 21.]
- K. Cha, K. W. Horch, and R. A. Normann. Mobility performance with a pixelized vision system. *Vision Res.*, 32(7): 1367–1372, 1992. [Cited on page 33.]
- G. J. Chader, J. Weiland, and M. S. Humayun. Artificial vision: needs, functioning, and testing of a retinal electronic prosthesis. *Prog. Brain Res.*, 175(09):317–332, 2009. [Cited on page 21.]
- O. Chapelle, T. Joachims, F. Radlinski, and Y. Yue. Large-scale validation and analysis of interleaved search evaluation. ACM Trans. Inf. Syst., 30(1), 2012. [Cited on page 42.]
- S. C. Chen, G. J. Suaning, J. W. Morley, and N. H. Lovell. Simulating prosthetic vision: I. Visual models of phosphenes. *Vision Res.*, 49(12):1493–1506, 2009a. [Cited on page 33.]
- S. C. Chen, G. J. Suaning, J. W. Morley, and N. H. Lovell. Rehabilitation regimes based upon psychophysical studies of prosthetic vision. *J. Neural Eng.*, 6(3), 2009b. [Cited on pages 20 and 33.]
- X. Chen, F. Wang, E. Fernandez, and P. R. Roelfsema. Shape perception via a high-channel-count neuroprosthesis in monkey visual cortex. *Science (80-.).*, 370(6521):1191–1196, 2020. [Cited on pages 15 and 66.]
- N. A. L. Chenais, M. J. I. Airaghi Leccardi, and D. Ghezzi. Photovoltaic retinal prosthesis restores high-resolution responses to single-pixel stimulation in blind retinas. *Commun. Mater.*, 2(1):1–16, 2021a. [Cited on pages 16 and 17.]
- N. A. L. Chenais, M. J. I. Airaghi Leccardi, and D. Ghezzi. Naturalistic spatiotemporal modulation of epiretinal stimulation increases the response persistence of retinal ganglion cell. *J. Neural Eng.*, 18(1), 2021b. [Cited on page 27.]
- J. S. Choi, A. J. Brockmeier, D. B. McNiel, L. M. von Kraus, J. C. Príncipe, and J. T. Francis. Eliciting naturalistic cortical responses with a sensory prosthesis via optimized microstimulation. *J. Neural Eng.*, 13(5):056007, 2016. [Cited on page 34.]

- W. Chu and Z. Ghahramani. Preference learning with Gaussian processes. In Proc. 22nd Int. Conf. Mach. Learn. -ICML '05, pages 137–144, New York, New York, USA, 2005. ACM Press. [Cited on pages 53 and 78.]
- B. G. Cleland, M. W. Dubin, and W. R. Levick. Sustained and transient neurones in the cat's retina and lateral geniculate nucleus. *J. Physiol.*, 217(2):473–496, 1971. [Cited on page 4.]
- R. J. Coffey. Deep brain stimulation devices: A brief technical history and review. *Artif. Organs*, 33(3):208–220, 2009. [Cited on page 13.]
- S. F. Cogan. Neural stimulation and recording electrodes. *Annu. Rev. Biomed. Eng.*, 10:275–309, 2008. [Cited on page 25.]
- D. Coppola and D. Purves. The extraordinarily rapid disappearance of entoptic images. *Proc. Natl. Acad. Sci. U. S. A.*, 93(15):8001–8004, 1996. [Cited on page 26.]
- M. Cox and B. De Vries. A parametric approach to Bayesian optimization with pairwise comparisons. In *Nips*, 2017. [Cited on pages 129 and 142.]
- S. I. Cunningham, J. D. Weiland, P. Bao, and B. S. Tjan. Visual cortex activation induced by tactile stimulation in late-blind individuals with retinitis pigmentosa. *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBS*, pages 2841–2844, 2011. [Cited on page 10.]
- C. A. Curcio and K. A. Allen. Topography of ganglion cells in human retina. *J. Comp. Neurol.*, 300(1):5–25, 1990. [Cited on pages 2, 2, 2, and 28.]
- C. A. Curcio, K. R. Sloan, R. E. Kalina, and A. E. Hendrickson. Human photoreceptor topography. J. Comp. Neurol., 292(4):497–523, 1990. [Cited on pages 1, 2, 2, 2, and 2.]
- C. A. Curcio, C. L. Millican, K. A. Allen, and R. E. Kalina. Aging of the human photoreceptor mosaic: Evidence for selective vulnerability of rods in central retina. *Investig. Ophthalmol. Vis. Sci.*, 34(12):3278–3296, 1993. [Cited on page 8.]
- L. da Cruz, F. K. Chen, A. Ahmado, J. Greenwood, and P. Coffey. RPE transplantation and its role in retinal disease. *Prog. Retin. Eye Res.*, 26(6):598–635, 2007. [Cited on page 12.]
- L. Da Cruz, B. F. Coley, J. Dorn, F. Merlini, E. Filley, P. Christopher, F. K. Chen, V. Wuyyuru, J. Sahel, P. Stanga, M. Humayun, R. J. Greenberg, and G. Dagnelie. The Argus II epiretinal prosthesis system allows letter and word reading and long-term function in patients with profound vision loss. *Br. J. Ophthalmol.*, 97(5):632–636, 2013. [Cited on pages 21, 22, 23, and 23.]
- L. da Cruz, J. D. Dorn, M. S. Humayun, G. Dagnelie, J. Handa, P. O. Barale, J. A. Sahel, P. E. Stanga, F. Hafezi, A. B. Safran, J. Salzmann, A. Santos, D. Birch, R. Spencer, A. V. Cideciyan, E. de Juan, J. L. Duncan, D. Eliott, A. Fawzi, L. C. Olmos de Koo, A. C. Ho, G. Brown, J. Haller, C. Regillo, L. V. Del Priore, A. Arditi, and R. J. Greenberg. Five-Year Safety and Performance Results from the Argus II Retinal Prosthesis System Clinical Trial. *Ophthalmology*, 123(10):2248–2254, 2016. [Cited on pages 22 and 48.]
- G. Dagnelie. Psychophysical evaluation for visual prosthesis. *Annu. Rev. Biomed. Eng.*, 10:339–368, 2008. [Cited on page 22.]
- A. C. Damianou and N. D. Lawrence. Deep Gaussian Processes. arXiv, 2013. [Cited on page 146.]
- P. M. Daniel and D. Whitteridge. The representation of the visual field on the cerebral cortex in monkeys. *J. Physiol.*, 159(2):203–221, 1961. [Cited on page 15.]
- R. Daschner, A. Rothermel, R. Rudorf, S. Rudorf, and A. Stett. Functionality and performance of the subretinal implant chip alpha AMS. *Sensors Mater.*, 30(2):179–192, 2018. [Cited on pages 24 and 24.]
- T. S. Davis, R. A. Parker, P. A. House, E. Bagley, S. Wendelken, R. A. Normann, and B. Greger. Spatial and temporal characteristics of V1 microstimulation during chronic implantation of a microelectrode array in a behaving macaque. *J. Neural Eng.*, 9(6), 2012. [Cited on pages 19, 25, and 25.]

- G. De Ath, R. M. Everson, A. A. M. Rahat, and J. E. Fieldsend. Greed is Good: Exploration and Exploitation Tradeoffs in Bayesian Optimisation. *arXiv*, 2019. [Cited on page 95.]
- G. De Ath, J. E. Fieldsend, and R. M. Everson. What do you Mean? The Role of the Mean Function in Bayesian Optimisation. *arXiv*, 2020. [Cited on page 142.]
- P. T. de Jong. Age-Related Macular Degeneration. N. Engl. J. Med., 355(14):1474–1485, 2006. [Cited on page 8.]
- A. De Palma, C. Mendler-Dünner, T. Parnell, A. Anghel, and H. Pozidis. Sampling Acquisition Functions for Batch Bayesian Optimization. *arXiv*, 2019. [Cited on page 147.]
- P. Degenaar, N. Grossman, M. A. Memon, J. Burrone, M. Dawson, E. Drakakis, M. Neil, and K. Nikolic. Optobionic vision A new genetically enhanced light on retinal prosthesis. *J. Neural Eng.*, 6(3), 2009. [Cited on pages 12 and 16.]
- M. P. Deisenroth, D. Fox, and C. E. Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(2):408–423, 2015. [Cited on page 108.]
- J. Delbeke. Electrodes and chronic optic nerve stimulation. *Biocybern. Biomed. Eng.*, 31(3):81–94, 2011. [Cited on page 14.]
- J. B. Demb and E. N. Pugh. Connexin36 Forms Synapses Essential for Night Vision. *Neuron*, 36(4):551–553, 2002. [Cited on page 5.]
- G. Denis, C. Jouffrais, C. Mailhes, and M. J. M. Mace. Simulated Prosthetic Vision: Improving text accessibility with retinal prostheses. 2014 36th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBC 2014, pages 1719–1722, 2014. [Cited on page 33.]
- S. Depeweg, J. M. Hernández-Lobato, F. Doshi-Velez, and S. Udluft. Decomposition of Uncertainty in Bayesian Deep Learning for Efficient and Risk-sensitive Learning. 35th Int. Conf. Mach. Learn. ICML 2018, 3:1920–1934, 2017. [Cited on pages 86 and 86.]
- O. Deroy and M. Auvray. Reading the world through the skin and ears: A new perspective on sensory substitution. *Front. Psychol.*, 3(NOV):1–13, 2012. [Cited on page 16.]
- I. Dewancker, M. McCourt, S. Clark, P. Hayes, A. Johnson, and G. Ke. A Stratified Analysis of Bayesian Optimization Methods. arXiv, 2016. [Cited on pages 100 and 124.]
- I. Dewancker, J. Bauer, and M. McCourt. Sequential Preference-Based Optimization. *arXiv*, 2017. [Cited on pages 51, 52, 75, 75, and 145.]
- J. DiGiovanna, B. Mahmoudi, J. Fortes, J. Principe, and J. Sanchez. Coadaptive Brain–Machine Interface via Reinforcement Learning. *IEEE Trans. Biomed. Eng.*, 56(1):54–64, 2009. [Cited on page 42.]
- Y. Ding, M. Kim, S. Kuindersma, and C. J. Walsh. Human-in-the-loop optimization of hip assistance with a soft exosuit during walking. *Sci. Robot.*, 3(15):1–9, 2018. [Cited on pages 41 and 72.]
- J. P. Dmochowski, A. Datta, M. Bikson, Y. Su, and L. C. Parra. Optimized multi-electrode stimulation increases focality and intensity at target. *J. Neural Eng.*, 8(4), 2011. [Cited on page 66.]
- W. H. Dobelle. Artificial vision for the blind by connecting a television camera to the visual cortex. *ASAIO J.*, 46(1): 3–9, 2000. [Cited on pages 15 and 20.]
- W. H. Dobelle and M. G. Mladejovsky. Phosphenes produced by electrical stimulation of human occipital cortex, and their application to the development of a prosthesis for the blind. *J. Physiol.*, 243(2):553–576, 1974. [Cited on page 28.]
- W. H. Dobelle, M. G. Mladejovsky, J. R. Evans, T. S. Roberts, and J. P. Girvin. 'Braille' reading by a blind volunteer by visual cortex stimulation. *Nature*, 259(5539):111–112, 1976. [Cited on page 20.]

- N. Donaldson and G. S. Brindley. The Historical Foundations of Bionics. In *Neurobionics Biomed. Eng. Neural Prostheses*, pages 1–37. John Wiley & Sons, Inc., Hoboken, NJ, USA, 2016. [Cited on page 15.]
- P. E. Donaldson. Experimental visual prosthesis. Proc. Inst. Electr. Eng., 120(2):281-299, 1973. [Cited on page 15.]
- J. D. Dorn, A. K. Ahuja, A. Caspi, L. Da Cruz, G. Dagnelie, J. A. Sahel, R. J. Greenberg, and M. J. McMahon. The detection of motion by blind subjects with the epiretinal 60-electrode (Argus II) retinal prosthesis. *JAMA Ophthalmol.*, 131(2):183–189, 2013. [Cited on pages 21, 21, 22, and 23.]
- M. Dudík, K. Hofmann, R. E. Schapire, A. Slivkins, and M. Zoghi. Contextual dueling bandits. *J. Mach. Learn. Res.*, 40(2015):1–25, 2015. [Cited on page 131.]
- J. Duebel, K. Marazova, and J. A. Sahel. Optogenetics. *Curr. Opin. Ophthalmol.*, 26(3):226–232, 2015. [Cited on page 12.]
- D. K. Duvenaud. *Automatic Model Construction with Gaussian Processes*. PhD thesis, University of Cambridge, 2014. [Cited on pages 82, 142, 142, and 151.]
- C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. *Proc. 10th Int. Conf. World Wide Web, WWW 2001,* pages 613–622, 2001. [Cited on pages 102 and 124.]
- R. Eckmiller. Learning retina implants with epiretinal contacts. *Ophthalmic Res.*, 29(5):281–289, 1997. [Cited on page 39.]
- R. Eckmiller, R. Hünermann, and M. Becker. Exploration of a dialog-based tunable retina encoder for retina implants. *Neurocomputing*, 26-27:1005–1011, 1999. [Cited on pages 20, 37, 39, 39, 66, and 136.]
- R. Eckmiller, D. Neumann, and O. Baruth. Tunable retina encoders for retina implants: why and how. *J. Neural Eng.*, 2(1):S91–S104, 2005. [Cited on pages 20, 23, 39, 39, and 66.]
- R. Eckmiller, R. Schatten, and O. Baruth. Portable Biomimetic Retina for Learning, Perception-based Image Acquisition. In 2007 Int. Jt. Conf. Neural Networks, pages 2436–2441. IEEE, 2007. [Cited on pages 39, 66, and 67.]
- T. L. Edwards, C. L. Cottriall, K. Xue, M. P. Simunovic, J. D. Ramsden, E. Zrenner, and R. E. MacLaren. Assessment of the Electronic Retinal Implant Alpha AMS in Restoring Vision to Blind Patients with End-Stage Retinitis Pigmentosa. *Ophthalmology*, 125(3):432–443, 2017. [Cited on pages 14, 23, and 25.]
- K. Eggensperger, F. Hutter, H. H. Hoos, and K. Leyton-Brown. Efficient benchmarking of hyperparameter optimizers via surrogates. *Proc. Natl. Conf. Artif. Intell.*, 2:1114–1120, 2015. [Cited on page 100.]
- G. Elachouri, I. Lee-Rivera, E. Clérin, M. Argentini, R. Fridlich, F. Blond, V. Ferracane, Y. Yang, W. Raffelsberger, J. Wan, J. Bennett, J. A. Sahel, D. J. Zack, and T. Léveillard. Thioredoxin rod-derived cone viability factor protects against photooxidative retinal damage. *Free Radic. Biol. Med.*, 81:22–29, 2015. [Cited on page 7.]
- M. T. M. Emmerich and A. H. Deutz. A tutorial on multiobjective optimization: fundamentals and evolutionary methods. *Nat. Comput.*, 17(3):585–609, 2018. [Cited on page 73.]
- C. Erickson-Davis and H. Korzybska. What do blind people "see" with retinal prostheses? Observations and qualitative reports of epiretinal implant users. *bioRxiv*, 2020. [Cited on pages 21, 22, 22, 47, 47, and 51.]
- T. Euler and R. H. Masland. Light-evoked responses of bipolar cells in a mammalian retina. *J. Neurophysiol.*, 83(4): 1817–1829, 2000. [Cited on page 3.]
- T. Euler, P. B. Detwiler, and W. Denk. Directionally selective calcium signals in dendrites of starburst amacrine cells. *Nature*, 418(6900):845–852, 2002. [Cited on page 4.]
- E. Famiglietti and H. Kolb. A bistratified amacrine cell and synaptic circuitry in the inner plexiform layer of the retina. *Brain Res.*, 84(2):293–300, 1975. [Cited on page 3.]

- V. H. Fan, L. E. Grosberg, S. S. Madugula, P. Hottowy, W. Dabrowski, A. Sher, A. M. Litke, and E. J. Chichilnisky. Epiretinal stimulation with local returns enhances selectivity at cellular resolution. *J. Neural Eng.*, 16(2), 2019. [Cited on page 28.]
- D. Feng and C. McCarthy. Enhancing scene structure in prosthetic vision using iso-disparity contour perturbance maps. *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBS*, pages 5283–5286, 2013. [Cited on page 46.]
- L. Ferlauto, M. J. I. Airaghi Leccardi, N. A. L. Chenais, S. C. A. Gilliéron, P. Vagni, M. Bevilacqua, T. J. Wolfensberger, K. Sivula, and D. Ghezzi. Design and validation of a foldable and photovoltaic wide-field epiretinal prosthesis. *Nat. Commun.*, 9(1):1–15, 2018. [Cited on pages 16, 26, and 27.]
- F. L. Ferris, C. P. Wilkinson, A. Bird, U. Chakravarthy, E. Chew, K. Csaky, and S. R. Sadda. Clinical classification of age-related macular degeneration. *Ophthalmology*, 120(4):844–851, 2013. [Cited on page 8.]
- I. Fine and G. M. Boynton. Pulse trains to percepts: The challenge of creating a perceptually intelligible world with sight recovery technologies. *Philos. Trans. R. Soc. B Biol. Sci.*, 370(1677), 2015. [Cited on pages 12, 24, 29, and 31.]
- I. Fine and R. A. Jacobs. Comparing perceptual learning tasks: A review. J. Vis., 2(2):190–203, 2002. [Cited on pages 10 and 10.]
- I. Fine, A. R. Wade, A. A. Brewer, M. G. May, D. F. Goodman, G. M. Boynton, B. A. Wandell, and D. I. MacLeod. Long-term deprivation affects visual perception and cortex. *Nat. Neurosci.*, 6(9):915–916, 2003. [Cited on pages 10, 10, and 10.]
- I. Fine, C. L. Cepko, and M. S. Landy. Vision research special issue: Sight restoration: Prosthetics, optogenetics and gene therapy. *Vision Res.*, 111:115–123, 2015. [Cited on page 12.]
- A. P. Finn, D. S. Grewal, and L. Vjzovic. Argus II retinal prosthesis system : a review of patient selection criteria , surgical considerations , and post-operative outcomes. *Clin. Ophthalmol.*, pages 1089–1097, 2018. [Cited on page 19.]
- K. M. Fitzgerald, G. W. Cibis, S. A. Giambrone, and D. J. Harris. Retinal signal transmission in Duchenne muscular dystrophy: Evidence for dysfunction in the photoreceptor/depolarizing bipolar cell pathway. *J. Clin. Invest.*, 93 (6):2425–2430, 1994. [Cited on page 31.]
- M. Fleckenstein, T. D. Keenan, R. H. Guymer, U. Chakravarthy, S. Schmitz-Valckenberg, C. C. Klaver, W. T. Wong, and E. Y. Chew. Age-related macular degeneration. *Nat. Rev. Dis. Prim.*, 7(1), 2021. [Cited on pages 8 and 9.]
- T. Flores, G. Goetz, X. Lei, and D. Palanker. Optimization of return electrodes in neurostimulating arrays. *J. Neural Eng.*, 13(3), 2016. [Cited on page 28.]
- O. Foerster. Beitrage zur pathophysiologie der sehbahn und der spehsphare. *J. Psychol. Neurol*, 39:463–485, 1929. [Cited on page 15.]
- A. I. J. Forrester, A. Sóbester, and A. J. Keane. *Engineering Design via Surrogate Modelling*. Wiley, 2008. [Cited on page 73.]
- P. Frazier, W. Powell, and S. Dayanik. The knowledge-gradient policy for correlated normal beliefs. *INFORMS J. Comput.*, 21(4):599–613, 2009. [Cited on pages 94, 121, and 121.]
- P. I. Frazier. A Tutorial on Bayesian Optimization. ArXiv, pages 1–22, 2018. [Cited on pages 51, 72, 84, and 91.]
- D. K. Freeman, D. K. Eddington, J. F. Rizzo, and S. I. Fried. Selective activation of neuronal targets with sinusoidal electric stimulation. J. Neurophysiol., 104(5):2778–2791, 2010. [Cited on page 30.]
- D. K. Freeman, J. F. Rizzo, and S. I. Fried. Encoding visual information in retinal ganglion cells with prosthetic stimulation. *J. Neural Eng.*, 8(3), 2011. [Cited on pages 27, 27, 27, and 68.]
- S. I. Fried, H. A. Hsueh, and F. S. Werblin. A method for generating precise temporal patterns of retinal spiking using prosthetic stimulation. *J. Neurophysiol.*, 95(2):970–978, 2006. [Cited on pages 27, 28, and 30.]

- S. I. Fried, C. Cai, and Q. Ren. High frequency electric stimulation of retinal neurons elicits physiological signaling patterns. *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBS*, pages 1077–1080, 2011. [Cited on page 30.]
- L. P. Fröhlich, E. D. Klenske, J. Vinogradska, C. Daniel, and M. N. Zeilinger. Noisy-Input Entropy Search for Efficient Robust Bayesian Optimization. *arXiv*, 2020. [Cited on page 73.]
- T. Fujikado, M. Kamei, H. Sakaguchi, H. Kanda, T. Endo, M. Hirota, T. Morimoto, K. Nishida, H. Kishima, Y. Terasawa, K. Oosawa, M. Ozawa, and K. Nishida. One-year outcome of 49-channel suprachoroidal–transretinal stimulation prosthesis in patients with advanced retinitis pigmentosa. *Investig. Ophthalmol. Vis. Sci.*, 57(14):6147– 6157, 2016. [Cited on pages 17 and 24.]
- C. Fülep, I. Kovács, K. Kránitz, and G. Erdei. Simulation of visual acuity by personalizable neuro-physiological model of the human eye. *Sci. Rep.*, 9(1):1–15, 2019. [Cited on page 127.]
- V. P. Gabel. Artificial Vision. Springer International Publishing, Cham, 2017. [Cited on page 23.]
- V. Gaillet, A. Cutrone, F. Artoni, P. Vagni, A. Mega Pratiwi, S. A. Romero, D. Lipucci Di Paola, S. Micera, and D. Ghezzi. Spatially selective activation of the visual cortex via intraneural stimulation of the optic nerve. *Nat. Biomed. Eng.*, 4(2):181–194, 2020. [Cited on page 14.]
- J. R. Gardner, K. Q. Weinberger, G. Malkomes, D. Barbour, R. Garnett, and J. P. Cunningham. Bayesian active model selection with an application to automated audiometry. *Adv. Neural Inf. Process. Syst.*, 2015-Janua:2386–2394, 2015. [Cited on page 128.]
- J. R. Gardner, C. Guo, K. Q. Weinberger, R. Garnett, and R. Grosse. Discovering and exploiting additive structure for Bayesian optimization. *Proc. 20th Int. Conf. Artif. Intell. Stat. AISTATS 2017*, 54, 2017. [Cited on pages 68, 139, and 142.]
- J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson. Gpytorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration. *Adv. Neural Inf. Process. Syst.*, 2018-Decem(NeurIPS):7576–7586, 2018. [Cited on page 113.]
- K. M. Gehrs. Complement, Age-Related Macular Degeneration and a Vision of the Future. *Arch. Ophthalmol.*, 128 (3):349, 2010. [Cited on page 6.]
- D. Ghezzi, M. R. Antognazza, M. Dal Maschio, E. Lanzarini, F. Benfenati, and G. Lanzani. A hybrid bioorganic interface for neuronal photoactivation. *Nat. Commun.*, 2(1), 2011. [Cited on page 16.]
- D. Ghezzi, M. R. Antognazza, R. MacCarone, S. Bellani, E. Lanzarini, N. Martino, M. Mete, G. Pertile, S. Bisti, G. Lanzani, and F. Benfenati. A polymer optoelectronic interface restores light sensitivity in blind rat retinas. *Nat. Photonics*, 7(5):400–406, 2013. [Cited on page 16.]
- D. Ginsbourger, R. L. Riche, and L. Carraro. A Multi-points Criterion for Deterministic Parallel Global Optimization based on Gaussian Processes. *HAL Prepr. hal00260579*, 1:1–30, 2008. [Cited on page 108.]
- G. Goetz, R. Smith, X. Lei, L. Galambos, T. Kamins, K. Mathieson, A. Sher, and D. Palanker. Contrast sensitivity with a subretinal prosthesis and implications for efficient delivery of visual information. *Investig. Ophthalmol. Vis. Sci.*, 56(12):7186–7194, 2015. [Cited on page 28.]
- J. R. Golden, C. Erickson-Davis, N. P. Cottaris, N. Parthasarathy, F. Rieke, D. H. Brainard, B. A. Wandell, and E. J. Chichilnisky. Simulation of visual perception and learning with a retinal prosthesis. *J. Neural Eng.*, 16(2):1–23, 2019. [Cited on pages 30, 30, 32, 33, 46, and 140.]
- F. Golemo, A. A. Taiga, A. Courville, and P.-Y. Oudeyer. Sim-to-Real Transfer with Neural-Augmented Robot Simulation. *Proc. 2nd Conf. Robot Learn.*, 87(CoRL):817–828, 2018. [Cited on page 143.]
- T. Gollisch and M. Meister. Eye Smarter than Scientists Believed: Neural Computations in Circuits of the Retina. *Neuron*, 65(2):150–164, 2010. [Cited on pages 2, 4, and 4.]

- J. González, M. Osborne, and N. D. Lawrence. GLASSES: Relieving the myopia of Bayesian optimisation. *Proc. 19th Int. Conf. Artif. Intell. Stat. AISTATS 2016*, pages 790–799, 2016. [Cited on page 146.]
- J. Gonzalez, Z. Dai, A. Damianou, and N. D. Lawrence. Preferential Bayesian optimization. 34th Int. Conf. Mach. Learn. ICML 2017, 3:2080–2089, 2017. [Cited on pages 51, 52, 75, 86, 87, 90, 90, 91, 97, 98, 98, 99, 100, 101, 101, 105, 108, 110, 111, 122, 123, and 160.]
- Y. S. Goo, K. N. Ahn, Y. J. Song, S. H. Ahn, S. K. Han, S. B. Ryu, and K. H. Kim. Spontaneous oscillatory rhythm in retinal activities of two retinal degeneration (rd1 and rd10) mice. *Korean J. Physiol. Pharmacol.*, 15(6):415–422, 2011. [Cited on page 10.]
- J. Görtler, R. Kehlbeck, and O. Deussen. A Visual Exploration of Gaussian Processes. *Distill*, 4(4), 2019. [Cited on page 75.]
- F. Gougoux, R. J. Zatorre, M. Lassonde, P. Voss, and F. Lepore. A functional neuroimaging study of sound localization: Visual cortex activity predicts performance in early-blind individuals. *PLoS Biol.*, 3(2):0324–0333, 2005. [Cited on page 10.]
- L. L. Grado, M. D. Johnson, and T. I. Netoff. Bayesian adaptive dual control of deep brain stimulation in a computational model of Parkinson's disease. *PLoS Comput. Biol.*, 14(12):1–23, 2018. [Cited on page 41.]
- R. J. Greenberg, T. J. Velte, M. S. Humayun, G. N. Scarlatis, and E. De Juan. A computational model of electrical stimulation of the retinal ganglion cell. *IEEE Trans. Biomed. Eng.*, 46(5):505–514, 1999. [Cited on pages 29 and 30.]
- S. H. Greenwald, A. Horsager, M. S. Humayun, R. J. Greenberg, M. J. McMahon, and I. Fine. Brightness as a Function of Current Amplitude in Human Retinal Electrical Stimulation. *Investig. Opthalmology Vis. Sci.*, 50(11):5017, 2009. [Cited on page 18.]
- W. M. Grill and J. T. Mortimer. Stimulus Waveforms for Selective Neural Stimulation. *IEEE Eng. Med. Biol. Mag.*, 14 (4):375–385, 1995. [Cited on page 13.]
- L. E. Grosberg, K. Ganesan, G. A. Goetz, S. S. Madugula, N. Bhaskhar, V. Fan, P. Li, P. Hottowy, W. Dabrowski, A. Sher, A. M. Litke, S. Mitra, and E. J. Chichilnisky. Activation of ganglion cells and axon bundles using epiretinal electrical stimulation. *J. Neurophysiol.*, 118(3):1457–1471, 2017. [Cited on pages 29, 30, 30, and 31.]
- S. Grover, G. A. Fishman, R. J. Anderson, M. S. Tozatti, J. R. Heckenlively, R. G. Weleber, A. O. Edwards, and J. Brown. Visual acuity impairment in patients with retinitis pigmentosa at age 45 years or older. *Ophthalmology*, 106(9):1780–1785, 1999. [Cited on page 6.]
- T. Guo, C. Y. Yang, D. Tsai, M. Muralidharan, G. J. Suaning, J. W. Morley, S. Dokos, and N. H. Lovell. Closed-loop efficient searching of optimal electrical stimulation parameters for preferential excitation of retinal ganglion cells. *Front. Neurosci.*, 12(MAR):1–12, 2018. [Cited on pages 31, 35, and 46.]
- P. R. Gupta and R. M. Huckfeldt. Gene therapy for inherited retinal degenerations: initial successes and future challenges. *J. Neural Eng.*, 14(5):051002, 2017. [Cited on page 11.]
- A. G. Habib, M. A. Cameron, G. J. Suaning, N. H. Lovell, and J. W. Morley. Spatially restricted electrical activation of retinal ganglion cells in the rabbit retina by hexapolar electrode return configuration. *J. Neural Eng.*, 10(3): 36013, 2013. [Cited on page 28.]
- A. E. Hadjinicolaou, H. Meffin, M. I. Maturana, S. L. Cloherty, and M. R. Ibbotson. Prosthetic vision: Devices, patient outcomes and retinal research. *Clin. Exp. Optom.*, 98(5):395–410, 2015. [Cited on pages 17 and 17.]
- Z. M. Hafed, K. Stingl, K. U. Bartz-Schmidt, F. Gekeler, and E. Zrenner. Oculomotor behavior of blind patients seeing with a subretinal visual implant. *Vision Res.*, 118:119–131, 2016. [Cited on pages 17, 26, 26, and 26.]
- M. Haim. The epidemiology of retinitis pigmentosa in Denmark. *Acta Ophthalmol. Scand.*, 80(2):1–34, 2002. [Cited on page 6.]

- L. R. Hallberg and A. Ringdahl. Living with cochlear implants: Experiences of 17 adult patients in Sweden. *Int. J. Audiol.*, 43(2):115–121, 2004. [Cited on page 11.]
- L. E. Hallum, G. Dagnelie, G. J. Suaning, and N. H. Lovell. Simulating auditory and visual sensorineural prostheses: A comparative review. *J. Neural Eng.*, 4(1), 2007. [Cited on page 33.]
- N. Han, S. Srivastava, A. Xu, D. Klein, and M. Beyeler. Deep Learning–Based Scene Simplification for Bionic Vision. *arXiv*, 1(1), 2021. [Cited on pages 20 and 33.]
- S. Hanneton, M. Auvray, and B. Durette. The Vibe: A versatile vision-to-audition sensory substitution device. *Appl. Bionics Biomech.*, 7(4):269–276, 2010. [Cited on page 16.]
- P. M. C. Harrison, R. Marjieh, F. Adolfi, P. van Rijn, M. Anglada-Tort, O. Tchernichovski, P. Larrouy-Maestri, and N. Jacoby. Gibbs Sampling with People. In *NeurIPS*, 2020. [Cited on page 145.]
- H. K. Hartline. the Response of Single Optic Nerve Fibers of the Vertebrate Eye To Illumination of the Retina. *Am. J. Physiol. Content*, 121(2):400–415, 1938. [Cited on page 4.]
- D. Hartong, E. Berson, and T. Dryja. Retinitis pigmentosa Prevalence and inheritance patterns. *Lancet*, 368:1795–1809, 2006. [Cited on pages 6, 6, and 7.]
- B. M. Harvey and S. O. Dumoulin. The relationship between cortical magnification factor and population receptive field size in human visual cortex: Constancies in cortical architecture. *J. Neurosci.*, 31(38):13604–13612, 2011. [Cited on page 15.]
- M. C. Hauer, P. J. Nasiatka, N. R. B. Stiles, J.-C. Lue, R. Agrawal, J. D. Weiland, M. S. Humayun, and A. R. Tanguay. Intraocular Camera for Retinal Prostheses: Optical Design. In *Front. Opt. 2007/Laser Sci. XXIII/Organic Mater. Devices Displays Energy Convers.*, page FThT1, Washington, D.C., 2007. OSA. [Cited on page 28.]
- R. Held, Y. Ostrovsky, B. Degelder, T. Gandhi, S. Ganesh, U. Mathur, and P. Sinha. The newly sighted fail to match seen with felt. *Nat. Neurosci.*, 14(5):551–553, 2011. [Cited on page 10.]
- F. Hell, C. Palleis, J. H. Mehrkens, T. Koeglsperger, and K. Bötzel. Deep brain stimulation programming 2.0: Future perspectives for target identification and adaptive closed loop stimulation. *Front. Neurol.*, 10(APR):1–11, 2019. [Cited on page 41.]
- P. Hennig and C. J. Schuler. Entropy search for information-efficient global optimization. *J. Mach. Learn. Res.*, 13: 1809–1837, 2012. [Cited on pages 89 and 90.]
- J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. *Adv. Neural Inf. Process. Syst.*, 1(January):918–926, 2014. [Cited on pages 99, 108, 108, 108, 108, 110, 111, 122, 132, 157, and 158.]
- P. Hickey and M. Stacy. Deep brain stimulation: A paradigm shifting approach to treat Parkinson's disease. *Front. Neurosci.*, 10(APR):1–11, 2016. [Cited on page 13.]
- C. H. Ho, R. J. Triolo, A. L. Elias, K. L. Kilgore, A. F. DiMarco, K. Bogie, A. H. Vette, M. L. Audu, R. Kobetic, S. R. Chang, K. M. Chan, S. Dukelow, D. J. Bourbeau, S. W. Brose, K. J. Gustafson, Z. H. Kiss, and V. K. Mushahwar. Functional electrical stimulation and Spinal Cord Injury. *Phys. Med. Rehabil. Clin. N. Am.*, 25(3):631–654, 2014. [Cited on page 13.]
- E. Ho, R. Smith, G. Goetz, I. Lei, L. Galambos, T. I. Kamins, J. Harris, K. Mathieson, D. Palanker, and A. Sher. Spatiotemporal characteristics of retinal response to network-mediated photovoltaic stimulation. *J. Neurophysiol.*, 119(2):389–400, 2018. [Cited on page 32.]
- L. R. Hochberg, D. Bacher, B. Jarosiewicz, N. Y. Masse, J. D. Simeral, J. Vogel, S. Haddadin, J. Liu, S. S. Cash, P. Van Der Smagt, and J. P. Donoghue. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375, 2012. [Cited on page 25.]

- B. A. Holden, T. R. Fricke, D. A. Wilson, M. Jong, K. S. Naidoo, P. Sankaridurg, T. Y. Wong, T. J. Naduvilath, and S. Resnikoff. Global Prevalence of Myopia and High Myopia and Temporal Trends from 2000 through 2050. *Ophthalmology*, 123(5):1036–1042, 2016. [Cited on page 127.]
- C. Hölig, J. Föcker, A. Best, B. Röder, and C. Büchel. Crossmodal plasticity in the fusiform gyrus of late blind individuals during voice recognition. *Neuroimage*, 103:374–382, 2014. [Cited on page 10.]
- S. Hombach, U. Janssen-Bienhold, G. Söhl, T. Schubert, H. Büssow, T. Ott, R. Weiler, and K. Willecke. Functional expression of connexin57 in horizontal cells of the mouse retina. *Eur. J. Neurosci.*, 19(10):2633–2640, 2004. [Cited on page 3.]
- D. C. Hood. Lower-level visual processing and models of light adaptation. *Annu. Rev. Psychol.*, 49:503–535, 1998. [Cited on page 2.]
- R. Hornig, T. Zehnder, and M. Velikay-Parel. The IMI retinal implant system. In Artif. Sight, pages 111–128. Springer, 2008. [Cited on pages 19, 20, and 23.]
- A. Horsager, S. H. Greenwald, J. D. Weiland, M. S. Humayun, R. J. Greenberg, M. J. McMahon, G. M. Boynton, and I. Fine. Predicting Visual Sensitivity in Retinal Prosthesis Patients. *Investig. Opthalmology Vis. Sci.*, 50(4):1483, 2009. [Cited on pages 18, 48, and 68.]
- A. Horsager, R. J. Greenberg, and I. Fine. Spatiotemporal interactions in retinal prosthesis subjects. *Investig. Oph-thalmol. Vis. Sci.*, 51(2):1223–1233, 2010. [Cited on page 48.]
- A. Horsager, G. M. Boynton, R. J. Greenberg, and I. Fine. Temporal interactions during paired-electrode stimulation in two retinal prosthesis subjects. *Investig. Ophthalmol. Vis. Sci.*, 52(1):549–557, 2011. [Cited on page 48.]
- J. C. Horton and D. R. Hocking. Monocular core zones and binocular border strips in primate striate cortex revealed by the contrasting effects of enucleation, eyelid suture, and retinal laser lesions on cytochrome oxidase activity. *J. Neurosci.*, 18(14):5433–5455, 1998. [Cited on pages 10 and 10.]
- J. C. Horton, M. Fahle, T. Mulder, and S. Trauzettel-Klosinski. Adaptation, perceptual learning, and plasticity of brain functions. *Graefe's Arch. Clin. Exp. Ophthalmol.*, 255(3):435–447, 2017. [Cited on page 10.]
- M. Hossein Shaker and E. Hüllermeier. Aleatoric and epistemic uncertainty with random forests. *arXiv*, pages 444–456, 2020. [Cited on page 86.]
- N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel. Bayesian Active Learning for Classification and Preference Learning. *arXiv*, pages 1–17, 2011. [Cited on pages 74, 87, 87, 89, 122, and 122.]
- N. Houlsby, J. M. Hernández-Lobato, F. Huszár, and Z. Ghahramani. Collaborative Gaussian processes for preference learning. *Adv. Neural Inf. Process. Syst.*, 3:2096–2104, 2012. [Cited on pages 53, 78, 78, 79, and 155.]
- D. Huang, T. T. Allen, W. I. Notz, and R. A. Miller. Sequential kriging optimization using multiple-fidelity evaluations. *Struct. Multidiscip. Optim.*, 32(5):369–382, 2006. [Cited on page 73.]
- E. Huber, J. M. Webster, A. A. Brewer, D. I. A. MacLeod, B. A. Wandell, G. M. Boynton, A. R. Wade, and I. Fine. A Lack of Experience-Dependent Plasticity After More Than a Decade of Recovered Sight. *Psychol. Sci.*, 26(4): 393–401, 2015. [Cited on page 10.]
- M. S. Humayun, M. Prince, E. De Juan, Y. Barron, M. Moskowitz, I. B. Klock, and A. H. Milam. Morphometric analysis of the extramacular retina from postmortem eyes with retinitis pigmentosa. *Investig. Ophthalmol. Vis. Sci.*, 40(1):143–148, 1999. [Cited on page 6.]
- M. S. Humayun, J. D. Dorn, L. da Cruz, G. Dagnelie, J.-A. Sahel, P. E. Stanga, A. V. Cideciyan, J. L. Duncan, D. Eliott, E. Filley, A. C. Ho, A. Santos, A. B. Safran, A. Arditi, L. V. Del Priore, and R. J. Greenberg. Interim Results from the International Trial of Second Sight's Visual Prosthesis. *Ophthalmology*, 119(4):779–788, 2012. [Cited on pages 21, 22, 23, 23, and 48.]

- F. Hutter, L. Kotthoff, and J. Vanschoren, editors. *Automated Machine Learning*. The Springer Series on Challenges in Machine Learning. Springer International Publishing, 2019. [Cited on page 139.]
- M. Im, P. Werginz, and S. I. Fried. Electric stimulus duration alters network-mediated responses depending on retinal ganglion cell type. *J. Neural Eng.*, 15(3):036010, 2018. [Cited on page 31.]
- J. Intoy and M. Rucci. Finely tuned eye movements enhance visual acuity. *Nat. Commun.*, 11(1):1–11, 2020. [Cited on page 26.]
- H. Jalali, I. Van Nieuwenhuyse, and V. Picheny. Comparison of Kriging-based algorithms for simulation optimization with heterogeneous noise. *Eur. J. Oper. Res.*, 261(1):279–301, 2017. [Cited on page 89.]
- N. M. Jansonius, J. Nevalainen, B. Selig, L. M. Zangwill, P. A. Sample, W. M. Budde, J. B. Jonas, W. A. Lagrèze, P. J. Airaksinen, R. Vonthein, L. A. Levin, J. Paetzold, and U. Schiefer. A mathematical description of nerve fiber bundle trajectories and their variability in the human retina. *Vision Res.*, 49(17):2157–2163, 2009. [Cited on pages 31, 50, 51, 54, 54, and 64.]
- M. Jasinska and E. Pyza. Circadian Plasticity of Mammalian Inhibitory Interneurons. *Neural Plast.*, 2017:1–12, 2017. [Cited on page 2.]
- K. Jemison. *Punishment and Prophecy: Blindness in Oedipus Rex, King Lear, and Jane Eyre*. PhD thesis, Florida Atlantic University, 2016. [Cited on page 1.]
- R. J. Jensen and J. F. Rizzo. Thresholds for activation of rabbit retinal ganglion cells with a subretinal electrode. *Exp. Eye Res.*, 83(2):367–373, 2006. [Cited on page 31.]
- R. J. Jensen and J. F. Rizzo. Responses of ganglion cells to repetitive electrical stimulation of the retina. *J. Neural Eng.*, 4(1), 2007. [Cited on page 27.]
- R. J. Jensen, J. F. Rizzo, O. R. Ziv, A. Grumet, and J. Wyatt. Thresholds for activation of rabbit retinal ganglion cells with an ultrafine, extracellular microelectrode. *Invest. Ophthalmol. Vis. Sci.*, 44(8):3533–3543, 2003. [Cited on pages 29 and 30.]
- L. H. Jepson, K. Hottowy Pawełand Mathieson, D. E. Gunning, W. Dabrowski, A. M. Litke, E. J. Chichilnisky, P. Hottowy, K. Mathieson, D. E. Gunning, W. Dabrowski, A. M. Litke, and E. J. Chichilnisky. Focal Electrical Stimulation of Major Ganglion Cell Types in the Primate Retina for the Design of Visual Prostheses. *J. Neurosci.*, 33(17):7194–7205, 2013. [Cited on page 32.]
- L. H. Jepson, P. Hottowy, K. Mathieson, D. E. Gunning, W. Dąbrowski, A. M. Litke, and E. J. Chichilnisky. Spatially Patterned Electrical Stimulation to Enhance Resolution of Retinal Prostheses. J. Neurosci., 34(14):4871–4881, 2014a. [Cited on pages 29 and 29.]
- L. H. Jepson, P. Hottowy, G. A. Weiner, W. Dabrowski, A. M. Litke, and E. J. Chichilnisky. High-fidelity reproduction of spatiotemporal visual signals for retinal prosthesis. *Neuron*, 83(1):87–92, 2014b. [Cited on pages 29 and 29.]
- S. Jiang, H. Chai, J. Gonzalez, and R. Garnett. BINOCULARS for Efficient, Nonmyopic Sequential Experimental Design. *arXiv*, 1(8):2–4, 2019. [Cited on pages 108, 108, 124, and 147.]
- S. Jiang, D. R. Jiang, M. Balandat, B. Karrer, J. R. Gardner, and R. Garnett. Efficient Nonmyopic Bayesian Optimization via One-Shot Multi-Step Trees. *arXiv*, 2020. [Cited on page 108.]
- B. W. Jones and R. E. Marc. Retinal remodeling during retinal degeneration. *Exp. Eye Res.*, 81(2):123–137, 2005. [Cited on page 8.]
- D. R. Jones, M. Schonlau, and W. J. Welch. Efficient Global Optimization of Expensive Black-Box Functions. J. Glob. Optim., 13(4):455–492, 1998. [Cited on pages 73 and 83.]
- P. R. Jones. QuestPlus: A MATLAB Implementation of the QUEST+ adaptive Psychometric Method. J. Open Res. Softw., 6, 2018. [Cited on pages 56, 62, and 69.]

- D. Kahneman and A. Tversky. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2):263, 1979. [Cited on pages 42, 52, and 75.]
- A. Kalia, L. A. Lesmes, M. Dorr, T. Gandhi, G. Chatterjee, S. Ganesh, P. J. Bex, and P. Sinha. Development of pattern vision following early and extended blindness. *Proc. Natl. Acad. Sci. U. S. A.*, 111(5):2035–2039, 2014. [Cited on page 10.]
- H. Kanda, T. Morimoto, T. Fujikado, Y. Tano, Y. Fukuda, and H. Sawai. Electrophysiological Studies of the Feasibility of Suprachoroidal-Transretinal Stimulation for Artificial Vision in Normal and RCS Rats. *Investig. Ophthalmol. Vis. Sci.*, 45(2):560–566, 2004. [Cited on pages 17 and 17.]
- K. Kandasamy, A. Krishnamurthy, J. Schneider, and B. Póczos. Parallelised bayesian optimisation via thompson sampling. *Int. Conf. Artif. Intell. Stat. AISTATS* 2018, 84:133–142, 2018. [Cited on page 144.]
- E. R. Kandel, J. H. Schwartz, and T. M. Jessell, editors. *Principles of Neural Science*. Elsevier, New York, third edition, 1991. [Cited on page 5.]
- J. Kasowski, N. Wu, and M. Beyeler. Towards Immersive Virtual Reality Simulations of Bionic Vision. *arXiv*, 1(1): 1–3, 2021. [Cited on page 33.]
- R. E. Kass and A. E. Raftery. Bayes Factors. J. Am. Stat. Assoc., 90(430):773, 1995. [Cited on pages 56 and 148.]
- K. Kauper, C. McGovern, S. Sherman, P. Heatherton, R. Rapoza, P. Stabila, B. Dean, A. Lee, S. Borges, B. Bouchard, and W. Tao. Two-year intraocular delivery of ciliary neurotrophic factor by encapsulated cell technology implants in patients with chronic retinal degenerative diseases. *Investig. Ophthalmol. Vis. Sci.*, 53(12):7484–7491, 2012. [Cited on page 11.]
- V. B. Kitiratschky, K. Stingl, B. Wilhelm, T. Peters, D. Besch, H. Sachs, F. Gekeler, K. U. Bartz-Schmidt, and E. Zrenner. Safety evaluation of "retina implant alpha IMS"—a prospective clinical trial. *Graefe's Arch. Clin. Exp. Ophthalmol.*, 253(3):381–387, 2015. [Cited on pages 17 and 25.]
- S. Klauke, M. Goertz, S. Rein, D. Hoehl, U. Thomas, R. Eckhorn, F. Bremmer, and T. Wachtler. Stimulation with a wireless intraocular epiretinal implant elicits visual percepts in blind humans. *Investig. Ophthalmol. Vis. Sci.*, 52 (1):449–455, 2011. [Cited on page 18.]
- A. Klein, S. Falkner, S. Bartels, P. Hennig, and F. Hutter. Fast Bayesian optimization of machine learning hyperparameters on large datasets. *Proc. 20th Int. Conf. Artif. Intell. Stat. AISTATS* 2017, 54, 2017. [Cited on page 131.]
- M. Kleiner, D. H. Brainard, D. G. Pelli, C. Broussard, T. Wolf, and D. Niehorster. What's new in Psychtoolbox-3? *Perception*, 36:S14, 2007. [Cited on page 69.]
- A. Kling, A. R. Gogliettino, N. P. Shah, E. G. Wu, N. Brackbill, A. Sher, A. M. Litke, R. A. Silva, and E. J. Chichilnisky. Functional Organization of Midget and Parasol Ganglion Cells in the Human Retina. *bioRxiv*, 019005:2020.08.07.240762, 2020. [Cited on page 4.]
- A. Koitschev, K. Stingl, K. U. Bartz-Schmidt, A. Braun, F. Gekeler, U. Greppmaier, H. Sachs, T. Peters, B. Wilhelm,
 E. Zrenner, and D. Besch. Extraocular surgical approach for placement of subretinal implants in blind patients: Lessons from cochlear-implants. J. Ophthalmol., 2015, 2015. [Cited on page 17.]
- K. Komeima, B. S. Rogers, L. Lu, and P. A. Campochiaro. Antioxidants reduce cone cell death in a model of retinitis pigmentosa. *Proc. Natl. Acad. Sci. U. S. A.*, 103(30):11300–11305, 2006. [Cited on pages 7 and 7.]
- K. Komeima, B. S. Rogers, and P. A. Campochiaro. Antioxidants slow photoreceptor cell death in mouse models of retinitis pigmentosa. J. Cell. Physiol., 213(3):809–15, 2007. [Cited on pages 7 and 7.]
- A. Krause and C. S. Ong. Contextual Gaussian process bandit optimization. Adv. Neural Inf. Process. Syst. 24 25th Annu. Conf. Neural Inf. Process. Syst. 2011, NIPS 2011, pages 1–9, 2011. [Cited on page 131.]
- F. Krause and H. Schum. *Neue Deutsche Chirurgie*. Kuttner, H., ed. KUTTNER, H., vol. 49a, pp. 482-486. Stuttgart: Enke., 1931. [Cited on page 15.]

- N. Kriegeskorte. Deep Neural Networks: A New Framework for Modeling Biological Vision and Brain Information Processing. *Annu. Rev. Vis. Sci.*, 1(1):417–446, 2015. [Cited on page 38.]
- N. Kumaran, A. T. Moore, R. G. Weleber, and M. Michaelides. Leber congenital amaurosis/early-onset severe retinal dystrophy: Clinical features, molecular genetics and therapeutic interventions. *Br. J. Ophthalmol.*, 101(9): 1147–1154, 2017. [Cited on page 12.]
- A. Kupcsik, D. Hsu, and W. S. Lee. Learning Dynamic Robot-to-Human Object Handover from Human Feedback. In *Springer Proc. Adv. Robot.*, volume 2, pages 161–176, 2018. [Cited on page 43.]
- M. J. Kyada, N. J. Killian, and J. S. Pezaris. Thalamic Visual Prosthesis Project. *Artif. Vis.*, pages 177–189, 2017. [Cited on page 15.]
- T. Z. Lauritzen, J. Harris, S. Mohand-Said, J. A. Sahel, J. D. Dorn, K. McClure, and R. J. Greenberg. Reading visual braille with a retinal prosthesis. *Biosyst. Biorobotics*, 6(November):29–39, 2014. [Cited on page 20.]
- A. H. Lay and A. K. Das. The role of neuromodulation in patients with neurogenic overactive bladder. *Curr. Urol. Rep.*, 13(5):343–347, 2012. [Cited on page 13.]
- M. Lázaro-Gredilla. Bayesian warped Gaussian processes. *Adv. Neural Inf. Process. Syst.*, 2:1619–1627, 2012. [Cited on page 146.]
- M. Lázaro-Gredilla, J. Quiñonero-Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal. Sparse spectrum gaussian process regression. J. Mach. Learn. Res., 11:1865–1881, 2010. [Cited on pages 108, 108, 109, and 157.]
- C. Le Roy. Sur quelques tentatives faites pour guérir diverses maladies par l'électricité. *Hist. l'Académie R. des Sci.,* pages 45–54, 1755. [Cited on page 13.]
- J. I. Lee and M. Im. Optimal Electric Stimulus Amplitude Improves the Selectivity between Responses of on Versus off Types of Retinal Ganglion Cells. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 27(10):2015–2024, 2019. [Cited on page 46.]
- S. W. Lee, D. K. Eddington, and S. I. Fried. Responses to pulsatile subretinal electric stimulation: Effects of amplitude and duration. *J. Neurophysiol.*, 109(7):1954–1968, 2013. [Cited on page 18.]
- S. Y. Lee, S. Usui, A. bakr Zafar, B. C. Oveson, Y. J. Jo, L. Lu, S. Masoudi, and P. A. Campochiaro. N-acetylcysteine promotes long-term survival of cones in a model of retinitis pigmentosa. *J. Cell. Physiol.*, 226(7):1843–1849, 2011. [Cited on page 7.]
- J. R. Lepird, M. P. Owen, and M. J. Kochenderfer. Bayesian preference elicitation for multiobjective engineering design optimization. *J. Aerosp. Inf. Syst.*, 12(10):634–645, 2015. [Cited on pages 42 and 138.]
- B. Letham, B. Karrer, G. Ottoni, and E. Bakshy. Constrained Bayesian optimization with noisy experiments. *Bayesian Anal.*, 14(2):495–519, 2019. [Cited on pages 73 and 91.]
- P. M. Lewis and J. V. Rosenfeld. Electrical stimulation of the brain and the development of cortical visual prostheses: An historical perspective. *Brain Res.*, 1630:208–224, 2016. [Cited on page 15.]
- C. Li, S. Gupta, S. Rana, V. Nguyen, S. Venkatesh, and A. Shilton. High dimensional Bayesian optimization using dropout. *IJCAI Int. Jt. Conf. Artif. Intell.*, 0:2096–2102, 2017. [Cited on pages 68 and 144.]
- L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *J. Mach. Learn. Res.*, 18:1–52, 2018. [Cited on page 74.]
- N. Lin, W. Fan, H. J. Sheedlo, J. E. Aschenbrenner, and J. E. Turner. Photoreceptor repair in response to RPE transplants in RCS rats: Outer segment regeneration. *Curr. Eye Res.*, 15(10):1069–1077, 1996. [Cited on page 12.]
- M. Lindauer, M. Feurer, K. Eggensperger, A. Biedenkapp, and F. Hutter. Towards Assessing the Impact of Bayesian Optimization's Own Hyperparameters. *arXiv*, 2019. [Cited on page 139.]

- J. Lindsey, S. A. Ocko, S. Ganguli, and S. Deny. A unified theory of early visual representations from retina to cortex through anatomically constrained deep cnNs. In *7th Int. Conf. Learn. Represent. ICLR 2019*, volume Part F1271, pages 417–425, 2019. [Cited on pages 38 and 38.]
- Y. Liu, N. R. B. Stiles, and M. Meister. Augmented Reality Powers a Cognitive Prosthesis for the Blind. *bioRxiv*, page 321265, 2018. [Cited on pages 16, 16, and 16.]
- H. Lorach, R. Benosman, O. Marre, S. H. Ieng, J. A. Sahel, and S. Picaud. Artificial retina: The multichannel processing of the mammalian retina achieved with a neuromorphic asynchronous light acquisition device. *J. Neural Eng.*, 9(6), 2012. [Cited on page 23.]
- H. Lorach, O. Marre, J.-a. Sahel, R. Benosman, and S. Picaud. Neural stimulation for visual rehabilitation: Advances and challenges. *J. Physiol.*, 107(5):421–431, 2013. [Cited on page 46.]
- H. Lorach, G. Goetz, R. Smith, X. Lei, Y. Mandel, T. Kamins, K. Mathieson, P. Huie, J. Harris, A. Sher, and D. Palanker. Photovoltaic restoration of sight with high visual acuity. *Nat. Med.*, 21(5):476–482, 2015. [Cited on pages 16, 17, 24, 24, 27, 29, and 32.]
- R. Lorenz. *Neuroadaptive Bayesian optimization: Implications for the Cognitive Sciences.* PhD thesis, Imperial College London, 2017. [Cited on page 96.]
- R. Lorenz, L. E. Simmons, R. P. Monti, J. L. Arthur, S. Limal, I. Laakso, R. Leech, and I. R. Violante. Efficiently searching through large tACS parameter spaces using closed-loop Bayesian optimization. *Brain Stimul.*, 12(6): 1484–1489, 2019. [Cited on pages 43 and 52.]
- J. D. Loudin, D. M. Simanovskii, K. Vijayraghavan, C. K. Sramek, A. F. Butterwick, P. Huie, G. Y. McLean, and D. V. Palanker. Optoelectronic retinal prosthesis: System design and performance. *J. Neural Eng.*, 4(1), 2007. [Cited on page 28.]
- K. H. Louie, M. N. Petrucci, L. L. Grado, C. Lu, P. J. Tuite, A. G. Lamperski, C. D. MacKinnon, S. E. Cooper, and T. I. Netoff. Semi-automated approaches to optimize deep brain stimulation parameters in Parkinson's disease. *J. Neuroeng. Rehabil.*, 18(1):83, 2021. [Cited on pages 42, 42, and 72.]
- L. Lü, M. Medo, C. H. Yeung, Y. C. Zhang, Z. K. Zhang, and T. Zhou. Recommender systems. *Phys. Rep.*, 519(1): 1–49, 2012. [Cited on page 40.]
- C. Lunghi, L. Galli-Resta, P. Binda, G. M. Cicchini, G. Placidi, B. Falsini, and M. C. Morrone. Visual cortical plasticity in retinitis pigmentosa. *Investig. Ophthalmol. Vis. Sci.*, 60(7):2753–2763, 2019. [Cited on pages 11 and 11.]
- Y. H. L. Luo and L. da Cruz. The Argus® II Retinal Prosthesis System. *Prog. Retin. Eye Res.*, 50:89–107, 2016. [Cited on pages 20 and 47.]
- R. E. MacLaren, R. A. Pearson, A. MacNeil, R. H. Douglas, T. E. Salt, M. Akimoto, A. Swaroop, J. C. Sowden, and R. R. Ali. Retinal repair by transplantation of photoreceptor precursors. *Nature*, 444(7116):203–207, 2006. [Cited on page 12.]
- M. A. MacNeil and R. H. Masland. Extreme diversity among amacrine cells: Implications for function. *Neuron*, 20 (5):971–982, 1998. [Cited on page 4.]
- S. Madugula, A. R. Gogliettino, M. Zaidi, G. Aggarwal, N. P. Shah, R. Vilkhu, M. Hays, H. Nguyen, R. A. Silva, E. J. Chichilnisky, A. Kling, N. P. Shah, R. Vilkhu, M. Hays, H. Nguyen, V. Fan, E. G. Wu, P. Hottowy, A. Sher, A. M. Litke, R. A. Silva, and E. J. Chichilnisky. Focal Electrical Stimulation of Human Retinal Ganglion Cells. *bioRxiv*, 019005:2020.08.23.263608, 2020. [Cited on pages 32 and 32.]
- G. Malkomes and R. Garnett. Automating Bayesian optimization with Bayesian optimization. *Adv. Neural Inf. Process. Syst.*, 2018-Decem(NeurIPS):5984–5994, 2018. [Cited on page 139.]
- M. Maniglia, V. Soler, B. Cottereau, and Y. Trotter. Spontaneous and training-induced cortical plasticity in MD patients: Hints from lateral masking. *Sci. Rep.*, 8(1):1–11, 2018. [Cited on pages 11 and 11.]

- R. E. Marc, B. W. Jones, C. B. Watt, and E. Strettoi. Neural remodeling in retinal degeneration. *Prog. Retin. Eye Res.*, 22(5):607–655, 2003. [Cited on page 8.]
- S. Martinez-Conde, S. L. Macknik, X. G. Troncoso, and T. A. Dyar. Microsaccades counteract visual fading during fixation. *Neuron*, 49(2):297–305, 2006. [Cited on page 26.]
- R. H. Masland. The Neuronal Organization of the Retina. Neuron, 76(2):266-280, 2012. [Cited on page 4.]
- Y. Masuda, S. O. Dumoulin, S. Nakadomari, and B. A. Wandell. V1 projection zone signals in human macular degeneration depend on task, not stimulus. *Cereb. Cortex*, 18(11):2483–2493, 2008. [Cited on pages 10 and 10.]
- Y. Masuda, H. Horiguchi, S. O. Dumoulin, A. Furuta, S. Miyauchi, S. Nakadomari, and B. A. Wandell. Taskdependent V1 responses in human retinitis pigmentosa. *Investig. Ophthalmol. Vis. Sci.*, 51(10):5356–5364, 2010. [Cited on page 10.]
- P. B. Matteucci, S. C. Chen, D. Tsai, C. W. D. Dodds, S. Dokos, J. W. Morley, N. H. Lovell, and G. J. Suaning. Current steering in retinal stimulation via a quasimonopolar stimulation paradigm. *Investig. Ophthalmol. Vis. Sci.*, 54(6): 4307–4320, 2013. [Cited on page 28.]
- M. I. Maturana, N. V. Apollo, A. E. Hadjinicolaou, D. J. Garrett, S. L. Cloherty, T. Kameneva, D. B. Grayden, M. R. Ibbotson, and H. Meffin. A Simple and Accurate Model to Predict Responses to Multi-electrode Stimulation in the Retina. *PLoS Comput. Biol.*, 12(4):1–26, 2016. [Cited on page 29.]
- M. I. Maturana, N. V. Apollo, D. J. Garrett, T. Kameneva, S. L. Cloherty, D. B. Grayden, A. N. Burkitt, M. R. Ibbotson, and H. Meffin. Electrical receptive fields of retinal ganglion cells: Influence of presynaptic neurons. *PLOS Comput. Biol.*, 14(2), 2018. [Cited on page 36.]
- J. F. Maya-Vetencourt, D. Ghezzi, M. R. Antognazza, E. Colombo, M. Mete, P. Feyen, A. Desii, A. Buschiazzo, M. Di Paolo, S. Di Marco, F. Ticconi, L. Emionite, D. Shmal, C. Marini, I. Donelli, G. Freddi, R. MacCarone, S. Bisti, G. Sambuceti, G. Pertile, G. Lanzani, and F. Benfenati. A fully organic retinal prosthesis restores vision in a rat model of degenerative blindness. *Nat. Mater.*, 16(6):681–689, 2017. [Cited on page 16.]
- R. W. McAllister. *Theology of blindness in the Hebrew scriptures*. PhD thesis, Andrews University, 2010. [Cited on page 1.]
- C. McCarthy, N. Barnes, and P. Lieby. Ground surface segmentation for navigation with a low resolution visual prosthesis. *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBS*, pages 4457–4460, 2011. [Cited on page 20.]
- M. E. McClements and R. E. MacLaren. Gene therapy for retinal disease. *Transl. Res.*, 161(4):241–254, 2013. [Cited on page 11.]
- D. McCreery, V. Pikov, and P. R. Troyk. Neuronal loss due to prolonged controlled-current stimulation with chronically implanted microelectrodes in the cat cerebral cortex. *J. Neural Eng.*, 7(3), 2010. [Cited on pages 25, 25, and 25.]
- D. B. McCreery, W. F. Agnew, T. G. Yuen, and L. Bullara. Charge density and charge per phase as cofactors in neural injury induced by electrical stimulation. *IEEE Trans. Biomed. Eng.*, 37(10):996–1001, 1990. [Cited on page 25.]
- C. C. McIntyre and W. M. Grill. Selective Microstimulation of Central Nervous System Neurons. *Ann. Biomed. Eng.*, 28(3):219–233, 2000. [Cited on page 13.]
- A. McKyton, I. Ben-Zion, R. Doron, and E. Zohary. The Limits of Shape Recognition following Late Emergence from Blindness. *Curr. Biol.*, 25(18):2373–2378, 2015. [Cited on page 10.]
- M. McLeod, M. A. Osborne, and S. J. Roberts. Practical Bayesian optimization for variable cost objectives. *arXiv*, 2017. [Cited on page 131.]
- N. E. Medeiros and C. A. Curcio. Preservation of ganglion cell layer neurons in age-related macular degeneration. *Investig. Ophthalmol. Vis. Sci.*, 42(3):795–803, 2001. [Cited on page 8.]

- P. Meijer. An experimental system for auditory image representations. *IEEE Trans. Biomed. Eng.*, 39(2):112–121, 1992. [Cited on page 16.]
- J. Menzler and G. Zeck. Network oscillations in rod-degenerated mouse retinas. J. Neurosci., 31(6):2280–2291, 2011. [Cited on page 10.]
- L. B. Merabet, J. F. Rizzo, A. Amedi, D. C. Somers, and A. Pascual-Leone. What blindness can tell us about seeing again: merging neuroplasticity and neuroprostheses. *Nat. Rev. Neurosci.*, 6(1):71–77, 2005. [Cited on page 8.]
- P. Mikkola, M. Todorović, J. Järvi, P. Rinke, and S. Kaski. Projective Preferential Bayesian Optimization. *arXiv*, 2020. [Cited on page 145.]
- P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002. [Cited on pages 153 and 153.]
- T. P. Minka. A family of algorithms for approximate Bayesian inference. PhD thesis, Massachusetts Institute of Technology, 2001. [Cited on pages 53 and 77.]
- R. M. Mirochnik and J. S. Pezaris. Contemporary approaches to visual prostheses. *Mil. Med. Res.*, 6(1):1–9, 2019. [Cited on pages 3 and 13.]
- J. Mockus. The Bayesian Approach to Local Optimization. Kluwer Academic Publishers, 1989. [Cited on page 83.]
- H. M. Mohammadi, E. Ghafar-Zadeh, and M. Sawan. An Image Processing Approach for Blind Mobility Facilitated Through Visual Intracortical Stimulation. *Artif. Organs*, 36(7):616–628, 2012. [Cited on page 20.]
- S. R. Montezuma, J. Loewenstein, C. Scholz, and J. F. Rizzo. Biocompatibility of materials implanted into the subretinal space of Yucatan pigs. *Investig. Ophthalmol. Vis. Sci.*, 47(8):3514–3522, 2006. [Cited on page 26.]
- J. K. Mueller and W. M. Grill. Model-based analysis of multiple electrode array stimulation for epiretinal visual prostheses. *J. Neural Eng.*, 10(3):036002, 2013. [Cited on pages 29 and 29.]
- S. Musa, D. R. Rand, D. J. Cott, J. Loo, C. Bartic, W. Eberle, B. Nuttin, and G. Borghs. Bottom-up SiO 2 embedded carbon nanotube electrodes with superior performance for integration in implantable neural microsystems. *ACS Nano*, 6(6):4615–4628, 2012. [Cited on page 25.]
- M. Mutný and A. Krause. Efficient high dimensional Bayesian optimization with additivity and quadrature fourier features. *Adv. Neural Inf. Process. Syst.*, 2018-Decem(NeurIPS):9005–9016, 2018. [Cited on pages 68, 109, and 142.]
- A. Najarpour-Foroushani, C. C. Pack, and M. Sawan. Cortical visual prostheses: From microstimulation to functional percept. J. Neural Eng., 15(2):021005, 2018. [Cited on page 15.]
- D. Nanduri, I. Fine, A. Horsager, G. M. Boynton, M. S. Humayun, R. J. Greenberg, and J. D. Weiland. Frequency and amplitude modulation have different effects on the percepts elicited by retinal stimulation. *Investig. Ophthalmol. Vis. Sci.*, 53(1):205–214, 2012. [Cited on pages 18, 18, 19, 48, and 69.]
- P. Nasiatka, A. Ahuja, N. R. B. Stiles, M. C. Hauer, R. N. Agrawal, R. Freda, D. Guven, M. S. Humayun, J. D. Weiland, and A. R. Tanguay. Intraocular Camera Design for Retinal Prostheses. In *Front. Opt.*, OSA Technical Digest Series, page FTh14, Washington, D.C., 2005. OSA. [Cited on page 28.]
- J. J. Nassi and E. M. Callaway. Parallel processing strategies of the primate visual system. *Nat. Rev. Neurosci.*, 10(5): 360–372, 2009. [Cited on pages 2 and 4.]
- S. Negi, R. Bhandari, L. Rieth, R. Van Wagenen, and F. Solzbacher. Neural electrode degradation from continuous electrical stimulation: Comparison of sputtered and activated iridium oxide. *J. Neurosci. Methods*, 186(1):8–17, 2010. [Cited on page 25.]
- S. Negi, R. Bhandari, and F. Solzbacher. A novel technique for increasing charge injection capacity of neural electrodes for efficacious and safe neural stimulation. *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBS*, pages 5142–5145, 2012. [Cited on pages 25 and 25.]

- D. C. Ng, C. E. Williams, P. J. Allen, S. Bai, C. S. Boyd, H. Meffin, M. E. Halpern, and E. Skafidas. Wireless power delivery for retinal prostheses. In *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBS*, pages 8356–8360. IEEE, 2011. [Cited on page 17.]
- J. S. Ng and S. C. Liem. Can the farnsworth D15 color vision test be defeated through practice? *Optom. Vis. Sci.*, 95 (5):452–456, 2018. [Cited on page 138.]
- Q. P. Nguyen, S. Tay, B. K. H. Low, and P. Jaillet. Top-\$k\$ Ranking Bayesian Optimization. *arXiv*, 2020. [Cited on pages 106, 145, and 145.]
- V. L. Nguyen, S. Destercke, and E. Hüllermeier. Epistemic Uncertainty Sampling. Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics), 11828 LNAI:72–86, 2019. [Cited on page 87.]
- J. Nielsen, J. Nielsen, and J. Larsen. Perception-based Personalization of Hearing Aids using Gaussian Processes and Active Learning. *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 23(1):1–1, 2014. [Cited on pages 52, 96, 96, 97, 101, 119, and 119.]
- J. B. Nielsen. *Systems for Personalization of Hearing Instruments*. PhD thesis, Technical University of Denmark, 2015. [Cited on pages 43 and 105.]
- S. Nirenberg and C. Pandarinath. Retinal prosthetic strategy with the capacity to restore normal vision. *Proc. Natl. Acad. Sci.*, 109(37):15012–15017, 2012. [Cited on pages 35, 46, 67, 136, and 136.]
- J. Oesterle, C. Behrens, C. Schröder, T. Herrmann, T. Euler, K. Franke, R. G. Smith, G. Zeck, and P. Berens. Bayesian inference for biophysical neuron models enables stimulus optimization for retinal neuroprosthetics. *Elife*, 9:1–38, 2020. [Cited on page 36.]
- Y. Ostrovsky, E. Meyers, S. Ganesh, U. Mathur, and P. Sinha. Visual parsing after recovery from blindness. *Psychol. Sci.*, 20(12):1484–1491, 2009. [Cited on page 10.]
- B. C. Oveson, T. Iwase, S. F. Hackett, S. Y. Lee, S. Usui, T. W. Sedlak, S. H. Snyder, P. A. Campochiaro, and J. U. Sung. Constituents of bile, bilirubin and TUDCA, protect against oxidative stress-induced retinal degeneration. *J. Neurochem.*, 116(1):144–153, 2011. [Cited on page 7.]
- P. Pacher, J. S. Beckman, and L. Liaudet. Nitric oxide and peroxynitrite in health and disease. *Physiol. Rev.*, 87(1): 315–424, 2007. [Cited on page 7.]
- L. Padnick-Silver, J. J. Kang Derwent, E. Giuliano, K. Narfström, and R. A. Linsenmeier. Retinal oxygenation and oxygen metabolism in abyssinian cats with a hereditary retinal degeneration. *Investig. Ophthalmol. Vis. Sci.*, 47(8): 3683–3689, 2006. [Cited on page 6.]
- D. Palanker, A. Vankov, P. Huie, and S. Baccus. Design of a high-resolution optoelectronic retinal prosthesis. *J. Neural Eng.*, 2(1), 2005. [Cited on page 28.]
- D. Palanker, Y. Le Mer, S. Mohand-Said, M. Muqit, and J. A. Sahel. Photovoltaic Restoration of Central Vision in Atrophic Age-Related Macular Degeneration. *Ophthalmology*, 127(8):1097–1104, 2020. [Cited on pages 24, 24, 26, and 46.]
- A. Paleyes, M. Pullin, M. Mahsereci, N. Lawrence, and J. Gonzalez. Emulation of physical processes with Emukit. In *NeurIPS*, 2019. [Cited on page 114.]
- M. G. Pandova, J. I. Al-Merjan, and N. A. Sadeq. Registered blindness in Kuwait–15 years of dynamic changes. *Ophthalmic Epidemiol.*, 26(2):75–83, 2019. [Cited on page 6.]
- F. Panetsos, A. Sanchez-Jimenez, E. D. D. Cerio, I. Diaz-Guemes, and F. M. Sanchez. Consistent phosphenes generated by electrical microstimulation of the visual thalamus. An experimental approach for thalamic visual neuroprostheses. *Front. Neurosci.*, 5(JUL):1–12, 2011. [Cited on pages 15, 15, and 15.]
- M. Parastarfeizabadi and A. Z. Kouzani. Advances in closed-loop deep brain stimulation devices. J. Neuroeng. Rehabil., 14(1):79, 2017. [Cited on page 41.]

- N. Parikh, L. Itti, M. Humayun, and J. Weiland. Performance of visually guided tasks using simulated prosthetic vision and saliency-based cues. *J. Neural Eng.*, 10(2), 2013. [Cited on page 20.]
- D. Pascolini and S. P. Mariotti. Global estimates of visual impairment: 2010. *Br. J. Ophthalmol.*, 96(5):614–618, 2012. [Cited on page 6.]
- J. K. Patel and C. B. Read. *Handbook of the Normal Distribution*. CRC Press, second edi edition, 1996. [Cited on pages 88 and 152.]
- M. Pearce, J. Klaise, and M. Groves. ConBO: Conditional Bayesian Optimization. *arXiv*, 2020. [Cited on pages 131, 131, and 132.]
- R. A. Pearson. Advances in repairing the degenerate retina by rod photoreceptor transplantation. *Biotechnol. Adv.*, 32(2):485–491, 2014. [Cited on page 12.]
- E. Peli. Testing Vision Is Not Testing For Vision. Transl. Vis. Sci. Technol., 9(13):32, 2020. [Cited on page 67.]
- D. G. Pelli. The VideoToolbox software for visual psychophysics: Transforming numbers into movies, 1997. [Cited on page 69.]
- A. Pérez Fornos, J. Sommerhalder, L. da Cruz, J. A. Sahel, S. Mohand-Said, F. Hafezi, and M. Pelizzone. Temporal Properties of Visual Perception on Electrical Stimulation of the Retina. *Investig. Opthalmology Vis. Sci.*, 53(6):2720, 2012. [Cited on pages 26, 26, and 68.]
- V. Perrone, H. Shen, M. Seeger, C. Archambeau, and R. Jenatton. Learning search spaces for Bayesian optimization: Another view of hyperparameter transfer learning. *NeurIPS*, 2019. [Cited on pages 66 and 142.]
- H. Petrs-Silva and R. Linden. Advances in gene therapy technologies to treat retinitis pigmentosa. *Clin. Ophthalmol.*, 8:127–136, 2013. [Cited on page 11.]
- J. S. Pezaris and E. N. Eskandar. Getting signals into the brain: Visual prosthetics through thalamic microstimulation. *Neurosurg. Focus*, 27(1), 2009. [Cited on page 15.]
- J. S. Pezaris and R. C. Reid. Demonstration of artificial visual percepts generated through thalamic microstimulation. *Proc. Natl. Acad. Sci. U. S. A.*, 104(18):7670–7675, 2007. [Cited on pages 14, 15, and 15.]
- R. L. Pfeiffer, R. E. Marc, and B. W. Jones. Persistent remodeling and neurodegeneration in late-stage retinal degeneration. *Prog. Retin. Eye Res.*, 74(July 2019):100771, 2020. [Cited on page 8.]
- P. M. Pilarski, M. R. Dawson, T. Degris, F. Fahimi, J. P. Carey, and R. S. Sutton. Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning. In *IEEE Int. Conf. Rehabil. Robot.*, pages 134–140. IEEE, 2011. [Cited on page 42.]
- R. Pinsler, P. Karkus, A. Kupcsik, D. Hsu, and W. S. Lee. Factored Contextual Policy Search with Bayesian Optimization. arXiv, pages 7242–7248, 2019. [Cited on page 131.]
- E. B. Plow, A. Pascual-Leone, and A. Machado. Brain Stimulation in the Treatment of Chronic Neuropathic and Non-Cancerous Pain. J. Pain, 13(5):411–424, 2012. [Cited on page 13.]
- V. S. Polikov, P. A. Tresco, and W. M. Reichert. Response of brain tissue to chronically implanted neural electrodes. J. Neurosci. Methods, 148(1):1–18, 2005. [Cited on page 25.]
- W. A. Pryor and G. L. Squadrito. The chemistry of peroxynitrite: A product from the reaction of nitric oxide with superoxide. *Am. J. Physiol. Lung Cell. Mol. Physiol.*, 268(5 12-5), 1995. [Cited on page 7.]
- T. W. Raasch. Spherocylindrical refractive errors and visual acuity. *Optom. Vis. Sci.*, 72(4):272–275, 1995. [Cited on page 128.]
- A. Rahimi and B. Recht. Random Features for Large-Scale Kernel Machines. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Adv. Neural Inf. Process. Syst.* 20, pages 1177–1184. Curran Associates, Inc., 2008. [Cited on pages 109 and 157.]

- A. Rai, R. Antonova, F. Meier, and C. G. Atkeson. Using simulation to improve sample-efficiency of Bayesian optimization for bipedal robots. *J. Mach. Learn. Res.*, 20:1–24, 2019. [Cited on page 143.]
- C. M. Ramsden, M. B. Powner, A. J. F. Carr, M. J. Smart, L. da Cruz, and P. J. Coffey. Stem cells in retinal regeneration: Past, present and future. *Dev.*, 140(12):2576–2585, 2013. [Cited on page 12.]
- P. V. Rao and L. L. Kupper. Ties in Paired-Comparison Experiments: A Generalization of the Bradley-Terry Model. J. Am. Stat. Assoc., 62(317):194–204, 1967. [Cited on page 145.]
- C. E. Rasmussen and K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006. [Cited on pages 52, 53, 54, 76, 76, 76, 77, 77, 136, and 161.]
- D. L. Rathbun, N. Ghorbani, H. Shabani, E. Zrenner, and Z. Hosseinzadeh. Spike-triggered average electrical stimuli as input filters for bionic vision A perspective. *J. Neural Eng.*, 15(6), 2018. [Cited on page 36.]
- F. Rattay and S. Resatz. Effective electrode configuration for selective stimulation with inner eye prostheses. *IEEE Trans. Biomed. Eng.*, 51(9):1659–1664, 2004. [Cited on page 29.]
- E. Raviola and R. F. Dacheux. Excitatory dyad synapse in rabbit retina. *Proc. Natl. Acad. Sci. U. S. A.*, 84(20):7324–7328, 1987. [Cited on page 3.]
- E. Richard, G. Goetz, and E. J. Chichilnisky. Recognizing retinal ganglion cells in the dark. *Adv. Neural Inf. Process. Syst.*, 2015-Janua:2476–2484, 2015. [Cited on page 32.]
- F. Rieke and D. A. Baylor. Single-photon detection by rod cells of the retina. *Rev. Mod. Phys.*, 70(3):1027–1036, 1998. [Cited on page 2.]
- L. A. Riggs, F. Ratliff, J. C. Cornsweet, and T. N. Cornsweet. The Disappearance of Steadily Fixated Visual Test Objects*. J. Opt. Soc. Am., 43(6):495, 1953. [Cited on page 26.]
- N. J. Rijkhoff, H. Wijkstra, P. E. Van Kerrebroeck, and F. M. Debruyne. Urinary bladder control by electrical stimulation: Review of electrical stimulation techniques in spinal cord injury. *Neurourol. Urodyn.*, 16(1):39–53, 1997. [Cited on page 13.]
- G. Riutort-Mayol, P. C. Burkner, M. R. Andersen, A. Solin, and A. Vehtari. Practical hilbert space approximate bayesian gaussian processes for probabilistic programming. *arXiv*, pages 1–33, 2020. [Cited on pages 110 and 158.]
- J. F. Rizzo. Update on retinal prosthetic research: The Boston retinal implant project. *J. Neuro-Ophthalmology*, 31(2): 160–168, 2011. [Cited on pages 17 and 48.]
- J. F. Rizzo, J. Wyatt, J. Loewenstein, S. Kelly, and D. Shire. Methods and Perceptual Thresholds for Short-Term Electrical Stimulation of Human Retina with Microelectrode Arrays. *Investig. Ophthalmol. Vis. Sci.*, 44(12):5355– 5361, 2003a. [Cited on pages 18 and 29.]
- J. F. Rizzo, J. Wyatt, J. Loewenstein, S. Kelly, and D. Shire. Perceptual Efficacy of Electrical Stimulation of Human Retina with a Microelectrode Array during Short-Term Surgical Trials. *Investig. Opthalmology Vis. Sci.*, 44(12): 5362, 2003b. [Cited on page 30.]
- S. Rizzo, C. Belting, L. Cinelli, L. Allegrini, F. Genovesi-Ebert, F. Barca, and E. Di Bartolo. The Argus II retinal prosthesis: 12-month outcomes from a single-study center. *Am. J. Ophthalmol.*, 157(6):1282–1290, 2014. [Cited on pages 21 and 21.]
- C. P. Robert. *Le choix bayésien*. Statistique et probabilités appliquées. Springer-Verlag, Paris, 2006. [Cited on pages 56 and 148.]
- R. W. Rodieck, K. F. Binmoeller, and J. Dineen. Parasol and midget ganglion cells of the human retina. J. Comp. Neurol., 233(1):115–132, 1985. [Cited on page 2.]

- P. Rolland, J. Scarlett, I. Bogunovic, and V. Cevher. High-dimensional Bayesian optimization via additive models with overlapping groups. *Int. Conf. Artif. Intell. Stat. AISTATS 2018*, 84:298–307, 2018. [Cited on pages 68 and 142.]
- S. Romeni, D. Zoccolan, and S. Micera. A machine learning framework to optimize optic nerve electrical stimulation for vision restoration. *Patterns*, page 100286, 2021. [Cited on pages 38 and 136.]
- M. G. Rosa, L. M. Schmid, and M. B. Calford. Responsiveness of cat area 17 after monocular inactivation: limitation of topographic plasticity in adult cortex. *J. Physiol.*, 482(3):589–608, 1995. [Cited on pages 10 and 10.]
- A. M. Rosse. Light and blindness in ancient Egypt. Acta med-hist Adriat, 8(2):221-238, 2010. [Cited on page 1.]
- P. J. Rousche and R. A. Normann. Chronic recording capability of the utah intracortical electrode array in cat sensory cortex. J. Neurosci. Methods, 82(1):1–15, 1998. [Cited on page 25.]
- D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen. A tutorial on Thompson sampling. *Found. Trends Mach. Learn.*, 11(1):1–96, 2018. [Cited on page 108.]
- H. P. Saal and S. J. Bensmaia. Biomimetic approaches to bionic touch through a peripheral nerve interface. *Neuropsychologia*, 79:344–353, 2015. [Cited on page 34.]
- N. Sabbah, C. N. Authié, N. Sanda, S. Mohand-Said, J. A. Sahel, and A. B. Safran. Importance of eye position on spatial localization in blind subjects wearing an argus II retinal prosthesis. *Investig. Ophthalmol. Vis. Sci.*, 55(12): 8259–8266, 2014. [Cited on page 28.]
- N. Sadato, A. Pascual-Leone, J. Grafman, V. Ibañez, M.-p. P. Deiber, G. Dold, and M. Hallett. Activation of the primary visual cortex by Braille reading in blind subjects. *Nature*, 380(6574):526–528, 1996. [Cited on page 10.]
- M. Saenz, L. B. Lewis, A. G. Huth, I. Fine, and C. Koch. Visual motion area MT + /V5 responds to auditory motion in human sight-recovery subjects. *J. Neurosci.*, 28(20):5141–5148, 2008. [Cited on page 10.]
- D. Sagi. Perceptual learning in Vision Research. Vision Res., 51(13):1552–1566, 2011. [Cited on page 21.]
- E. Sagi and M. A. Svirsky. Deactivating cochlear implant electrodes to improve speech perception: A computational approach. *Hear. Res.*, 370:316–328, 2018. [Cited on page 30.]
- J.-a. Sahel, E. Boulanger-Scemama, C. Pagot, A. Arleo, F. Galluppi, J. N. Martel, S. D. Esposti, A. Delaux, J.-B. de Saint Aubert, C. de Montleau, E. Gutman, I. Audo, J. Duebel, S. Picaud, D. Dalkara, L. Blouin, M. Taiel, and B. Roska. Partial recovery of visual function in a blind patient after optogenetic therapy. *Nat. Med.*, 2021. [Cited on page 12.]
- H. Sakaguchi, M. Kamei, T. Fujikado, E. Yonezawa, M. Ozawa, C. Cecilia-Gonzalez, O. Ustariz-Gonzalez, H. Quiroz-Mercado, and Y. Tano. Artificial vision by direct optic nerve electrode (AV-DONE) implantation in a blind patient with retinitis pigmentosa. J. Artif. Organs, 12(3):206–209, 2009. [Cited on page 14.]
- J. R. Sanes and R. H. Masland. The Types of Retinal Ganglion Cells: Current Status and Implications for Neuronal Classification. *Annu. Rev. Neurosci.*, 38:221–246, 2015. [Cited on page 31.]
- A. Santos. Preservation of the Inner Retina in Retinitis Pigmentosa. *Arch. Ophthalmol.*, 115(4):511, 1997. [Cited on page 6.]
- T. Santos-Ferreira, K. Postel, H. Stutzki, T. Kurth, G. Zeck, and M. Ader. Daylight vision repair by cell transplantation. *Stem Cells*, 33(1):79–90, 2015. [Cited on page 12.]
- M. M. Sanz, L. E. Johnson, S. Ahuja, P. A. Ekström, J. Romero, and T. van Veen. Significant photoreceptor rescue by treatment with a combination of antioxidants in an animal model for retinal degeneration. *Neuroscience*, 145(3): 1120–1129, 2007. [Cited on page 7.]
- S. D. Schwartz, J. P. Hubschman, G. Heilwell, V. Franco-Cardenas, C. K. Pan, R. M. Ostrick, E. Mickunas, R. Gay, I. Klimanskaya, and R. Lanza. Embryonic stem cell trials for macular degeneration: A preliminary report. *Lancet*, 379(9817):713–720, 2012. [Cited on page 12.]

- J. M. Seddon. Macular degeneration epidemiology: Nature-nurture, lifestyle factors, genetic risk, and geneenvironment interactions – the weisenfeld award lecture. *Investig. Ophthalmol. Vis. Sci.*, 58(14):6513–6528, 2017. [Cited on page 8.]
- M. Seeger. Notes on Minka's expectation propagation for Gaussian process classification, 2002. [Cited on pages 53 and 77.]
- C. Sekirnjak, P. Hottowy, A. Sher, W. Dabrowski, A. M. Litke, and E. J. Chichilnisky. Electrical Stimulation of Mammalian Retinal Ganglion Cells With Multielectrode Arrays. J. Neurophysiol., 95(6):3311–3327, 2006. [Cited on pages 17, 27, and 28.]
- C. Sekirnjak, P. Hottowy, A. Sher, W. Dabrowski, A. M. Litke, and E. J. Chichilnisky. High-Resolution Electrical Stimulation of Primate Retina for Epiretinal Implant Design. *J. Neurosci.*, 28(17):4446–4456, 2008. [Cited on pages 28 and 32.]
- N. B. Shah, S. Balakrishnan, J. Bradley, A. Parekh, K. Ramchandran, and M. Wainwright. When is it Better to Compare than to Score? *arXiv*, pages 1–21, 2014. [Cited on pages 52 and 75.]
- N. P. Shah and E. J. Chichilnisky. Computational challenges and opportunities for a bi-directional artificial retina. *J. Neural Eng.*, 17(5):055002, 2020. [Cited on page 32.]
- N. P. Shah, S. Madugula, E. J. Chichilnisky, J. Shlens, Y. Singer, and G. Brain. Learning a neural response metric for retinal prosthesis. *bioRxiv*, pages 1–13, 2017. [Cited on pages 38 and 67.]
- N. P. Shah, S. Madugula, L. Grosberg, G. Mena, P. Tandon, P. Hottowy, A. Sher, A. Litke, S. Mitra, and E. J. Chichilnisky. Optimization of Electrical Stimulation for a High-Fidelity Artificial Retina. In *Int. IEEE/EMBS Conf. Neural Eng. NER*, volume 2019-March, pages 714–718. IEEE, 2019. [Cited on pages 36, 46, 46, 66, and 136.]
- B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE*, 104(1):148–175, 2016. [Cited on pages 51 and 72.]
- R. V. Shannon. Advances in auditory prostheses. Curr. Opin. Neurol., 25(1):61-66, 2012. [Cited on pages 11 and 30.]
- Y. Shao, G. A. Keliris, A. Papanikolaou, M. D. Fischer, D. Zobor, H. Jägle, N. K. Logothetis, and S. M. Smirnakis. Visual cortex organisation in a macaque monkey with macular degeneration. *Eur. J. Neurosci.*, 38(10):3456–3464, 2013. [Cited on page 10.]
- R. K. Shepherd, M. N. Shivdasani, D. A. X. Nayagam, C. E. Williams, and P. J. Blamey. Visual prostheses for the blind. *Trends Biotechnol.*, 31(10):562–571, 2013. [Cited on page 13.]
- P. A. Sieving, R. C. Caruso, W. Tao, H. R. Coleman, D. J. Thompson, K. R. Fullmer, and R. A. Bush. Ciliary neurotrophic factor (CNTF) for human retinal degeneration: Phase I trial of CNTF delivered by encapsulated cell intraocular implants. *Proc. Natl. Acad. Sci. U. S. A.*, 103(10):3896–3901, 2006. [Cited on page 11.]
- E. Siivola, A. K. Dhaka, M. R. Andersen, J. Gonzalez, P. G. Moreno, and A. Vehtari. Preferential Batch Bayesian Optimization. *arXiv*, 2020. [Cited on pages 105, 106, 108, 115, 145, and 145.]
- R. Šikl, M. Šimeček, M. Porubanová-Norquist, O. Bezdíček, J. Kremláček, P. Stodůlka, I. Fine, and Y. Ostrovsky. Vision after 53 years of blindness. *Iperception.*, 4(8):498–507, 2013. [Cited on pages 10, 10, 10, and 10.]
- E. P. Simoncelli, H. R. Sheikh, A. C. Bovik, and Z. Wang. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. image Process.*, 13(4):600–612, 2004. [Cited on pages 51 and 143.]
- N. C. Sinclair, M. N. Shivdasani, T. Perera, L. N. Gillespie, H. J. McDermott, L. N. Ayton, and P. J. Blamey. The appearance of phosphenes elicited using a suprachoroidal retinal prosthesis. *Investig. Ophthalmol. Vis. Sci.*, 57 (11):4948–4961, 2016. [Cited on pages 13 and 30.]
- P. Sinha, G. Chatterjee, T. Gandhi, and A. Kalia. Restoring Vision through "Project Prakash": The Opportunities for Merging Science and Service. *PLoS Biol.*, 11(12), 2013. [Cited on pages 10 and 10.]

- S. M. Smirnakis, A. A. Brewer, M. C. Schmid, A. S. Tolias, A. Schüz, M. Augath, W. Inhoffen, B. A. Wandell, and N. K. Logothetis. Lack of long-term cortical reorganization after macaque retinal lesions. *Nature*, 435(7040):300–307, 2005. [Cited on pages 10 and 10.]
- J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian Optimization of Machine Learning Algorithms. In *Proc.* 25th Int. Conf. Neural Inf. Process. Syst., volume 2, pages 2951–2959, 2012. [Cited on pages 54, 73, 139, and 146.]
- J. Snoek, K. Swersky, R. Zemel, and R. P. Adams. Input warping for Bayesian optimization of non-stationary functions. *31st Int. Conf. Mach. Learn. ICML* 2014, 5:3654–3662, 2014. [Cited on page 146.]
- J. Snoek, O. Ripped, K. Swersky, R. Kiros, N. Satish, N. Sundaram, M. M. A. Patwary, Prabhat, and R. P. Adams. Scalable Bayesian optimization using deep neural networks. *32nd Int. Conf. Mach. Learn. ICML* 2015, 3:2161–2170, 2015. [Cited on page 76.]
- A. Solin and S. Särkkä. Hilbert space methods for reduced-rank Gaussian process regression. *Stat. Comput.*, 30(2): 419–446, 2020. [Cited on pages 110, 110, 111, 113, 122, and 158.]
- X. D. Song, R. Garnett, and D. L. Barbour. Psychometric function estimation by probabilistic classification. J. Acoust. Soc. Am., 141(4):2513–2525, 2017. [Cited on page 128.]
- X. D. Song, K. A. Sukesan, and D. L. Barbour. Bayesian active probabilistic classification for psychometric field estimation. *Attention, Perception, Psychophys.*, 80(3):798–812, 2018. [Cited on page 128.]
- M. J. Spencer, T. Kameneva, D. B. Grayden, H. Meffin, and A. N. Burkitt. Global activity shaping strategies for a retinal implant. *J. Neural Eng.*, 16(2):026008, 2018. [Cited on pages 29, 46, and 66.]
- J. T. Springenberg, A. Klein, S. Falkner, and F. Hutter. Bayesian optimization with Robust Bayesian neural networks. In *Adv. Neural Inf. Process. Syst.*, pages 4141–4149, 2016. [Cited on page 76.]
- N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *ICML 2010 - Proceedings*, 27th Int. Conf. Mach. Learn., pages 1015–1022, 2010. [Cited on pages 115, 122, and 146.]
- P. Stanga, J. Sahel, S. Mohand-Said, L. DaCruz, A. Caspi, F. Merlini, R. Greenberg, and A. I. I. S. Group. Face Detection using the Argus® II Retinal Prosthesis System. *Invest. Ophthalmol. Vis. Sci.*, 54(15):1766, 2013. [Cited on page 20.]
- A. Stett, W. Barth, S. Weiss, H. Haemmerle, and E. Zrenner. Electrical multisite stimulation of the isolated chicken retina. *Vision Res.*, 40(13):1785–1795, 2000. [Cited on pages 16 and 17.]
- G. A. Stevens, R. A. White, S. R. Flaxman, H. Price, J. B. Jonas, J. Keeffe, J. Leasher, K. Naidoo, K. Pesudovs, S. Resnikoff, H. Taylor, and R. R. Bourne. Global prevalence of vision impairment and blindness: Magnitude and temporal trends, 1990-2010. *Ophthalmology*, 120(12):2377–2384, 2013. [Cited on page 6.]
- K. Stingl, K. U. Bartz-Schmidt, D. Besch, C. K. Chee, C. L. Cottriall, F. Gekeler, M. Groppe, T. L. Jackson, R. E. MacLaren, A. Koitschev, A. Kusnyerik, J. Neffendorf, J. Nemeth, M. A. N. Naeem, T. Peters, J. D. Ramsden, H. Sachs, A. Simpson, M. S. Singh, B. Wilhelm, D. Wong, and E. Zrenner. Subretinal Visual Implant Alpha IMS Clinical trial interim report. *Vis. Res.*, 111:149–160, 2015. [Cited on pages 18, 26, and 26.]
- K. Stingl, R. Schippert, K. U. Bartz-Schmidt, D. Besch, C. L. Cottriall, T. L. Edwards, F. Gekeler, U. Greppmaier, K. Kiel, A. Koitschev, L. Kühlewein, R. E. MacLaren, J. D. Ramsden, J. Roider, A. Rothermel, H. Sachs, G. S. Schröder, J. Tode, N. Troelenberg, and E. Zrenner. Interim results of a multicenter trial with the new electronic subretinal implant alpha AMS in 15 patients blind from inherited retinal degenerations. *Front. Neurosci.*, 11 (AUG):445, 2017. [Cited on pages 18, 23, 24, and 46.]
- K. T. K. Stingl, K. U. Bartz-Schmidt, D. Besch, A. Braun, A. Bruckmann, F. Gekeler, U. Greppmaier, S. Hipp, G. Hortdorfer, C. Kernstock, A. Koitschev, A. Kusnyerik, H. Sachs, A. Schatz, K. T. K. Stingl, T. Peters, B. Wilhelm, and E. Zrenner. Artificial vision with wirelessly powered subretinal electronic implant alpha-IMS. *Proc. R. Soc. B Biol. Sci.*, 280(1757):1–8, 2013. [Cited on pages 21, 22, 23, and 39.]

- B. W. Streeten. The Sudanophilic Granules of the Human Retinal Pigment Epithelium. *Arch. Ophthalmol.*, 66(3): 391–398, 1961. [Cited on page 5.]
- E. Strettoi. A survey of retinal remodeling. Front. Cell. Neurosci., 9(DEC):1-4, 2015. [Cited on page 8.]
- Y. Sui and J. W. Burdick. Correlational dueling bandits with application to clinical treatment in large decision spaces. *IJCAI Int. Jt. Conf. Artif. Intell.*, pages 2793–2799, 2017. [Cited on pages 101, 105, 105, 106, 107, 115, 123, 145, and 145.]
- Y. Sui, V. Zhuang, J. W. Burdick, and Y. Yue. Multi-dueling bandits with dependent arms. *Uncertain. Artif. Intell. -Proc. 33rd Conf. UAI 2017*, 2017. [Cited on pages 99, 108, and 122.]
- C.-H. Sung and J.-Z. Chuang. The cell biology of vision. J. Cell Biol., 190(6):953-963, 2010. [Cited on page 3.]
- S. Surjanovic and D. Bingham. Virtual Library of Simulation Experiments: Test Functions and Datasets. Retrieved July 29, 2021, from \url{http://www.sfu.ca/~ssurjano}, 2021. [Cited on pages 100, 102, and 123.]
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, second edi edition, 2018. [Cited on page 42.]
- S. Tafazoli, C. J. MacDowell, Z. Che, K. C. Letai, C. Steinhardt, and T. J. Buschman. Learning to control the brain through adaptive closed-loop patterned stimulation. *bioRxiv*, pages 1–41, 2020. [Cited on pages 35 and 38.]
- P. Tandon, N. Bhaskhar, N. Shah, S. Madugula, L. Grosberg, V. H. Fan, P. Hottowy, A. Sher, E. J. Chichilnisky, and S. Mitra. Automatic Identification and Avoidance of Axon Bundle Activation for Epiretinal Prosthesis. *bioRxiv*, 021271:1–12, 2021. [Cited on page 30.]
- E. J. Tehovnik. Electrical stimulation of neural tissue to evoke behavioral responses. *J. Neurosci. Methods*, 65(1):1–17, 1996. [Cited on pages 13 and 28.]
- E. J. Tehovnik, A. S. Tolias, F. Sultan, W. M. Slocum, and N. K. Logothetis. Direct and indirect activation of cortical neurons by electrical microstimulation. J. Neurophysiol., 96(2):512–521, 2006. [Cited on page 13.]
- M. Tesch, J. Schneider, and H. Choset. Expensive function optimization with stochastic binary outcomes. 30th Int. Conf. Mach. Learn. ICML 2013, 28(PART 3):2320–2328, 2013. [Cited on pages 74, 74, 90, 91, 92, 95, 95, 103, 104, 104, 119, 120, and 122.]
- N. Thatte, H. Duan, and H. Geyer. A Sample-Efficient Black-Box Optimizer to Train Policies for Human-in-the-Loop Systems with User Preferences. *IEEE Robot. Autom. Lett.*, 2(2):993–1000, 2017. [Cited on pages 99, 101, and 132.]
- G. E. Thomas. Remark ASR 65: A Remark on Algorithm AS76: An Integral Useful in Calculating Non-Central t and Bivariate Normal Probabilities. *Appl. Stat.*, 35(3):310, 1986. [Cited on pages 53 and 88.]
- P. Thomas, M. Branicky, A. van den Bogert, and K. Jagodnik. Application of the Actor-Critic Architecture to Functional Electrical Stimulation Control of a Human Arm. Proc. ... Innov. Appl. Artif. Intell. Conf. Innov. Appl. Artif. Intell. Conf., 2009(1):165–172, 2009. [Cited on page 42.]
- W. R. Thompson. On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*, 25(3/4):285, 1933. [Cited on pages 61, 94, 108, and 122.]
- S. A. Titchener, M. N. Shivdasani, J. B. Fallon, and M. A. Petoe. Gaze compensation as a technique for improving hand– eye coordination in prosthetic vision. *Transl. Vis. Sci. Technol.*, 7(1):11–13, 2018. [Cited on page 28.]
- T. Tokutake, N. Mita, K.-I. Kawamoto, K. Kani, and A. Tabuchi. Relation between Visual Acuity and Slope of Psychometric Function in Young Adults. *Iperception.*, 2(4):308–308, 2011. [Cited on page 127.]
- S. Toscano-Palmerin and P. I. Frazier. Bayesian Optimization with Expensive Integrands. *arXiv*, pages 1–52, 2018. [Cited on pages 131 and 131.]

- D. Tsai, S. Chen, D. A. Protti, J. W. Morley, G. J. Suaning, and N. H. Lovell. Responses of Retinal Ganglion Cells to Extracellular Electrical Stimulation, from Single Cell to Population: Model-Based Analysis. *PLoS One*, 7(12), 2012. [Cited on page 29.]
- M. Tucker, E. Novoseller, C. Kann, Y. Sui, Y. Yue, J. Burdick, and A. D. Ames. Preference-Based Learning for Exoskeleton Gait Optimization. *arXiv*, 2019. [Cited on page 42.]
- U. Tulunay-Keesey. Fading of Stabilized Retinal Images. J. Opt. Soc. Am., 72(4):440-447, 1982. [Cited on page 26.]
- R. Turner, D. Eriksson, M. McCourt, J. Kiili, E. Laaksonen, Z. Xu, and I. Guyon. Bayesian Optimization is Superior to Random Search for Machine Learning Hyperparameter Tuning: Analysis of the Black-Box Optimization Challenge 2020. *arXiv*, 2021. [Cited on pages 72, 135, and 146.]
- P. Twyford, C. Cai, and S. Fried. Differential responses to high-frequency electrical stimulation in on and off retinal ganglion cells. *J. Neural Eng.*, 11(2), 2014. [Cited on pages 32 and 46.]
- S. Usui, K. Komeima, S. Y. Lee, Y. J. Jo, S. Ueno, B. S. Rogers, Z. Wu, J. Shen, L. Lu, B. C. Oveson, P. S. Rabinovitch, and P. A. Campochiaro. Increased expression of catalase and superoxide dismutase 2 reduces cone cell death in retinitis pigmentosa. *Mol. Ther.*, 17(5):778–786, 2009. [Cited on page 7.]
- S. Usui, B. C. Oveson, T. Iwase, L. Lu, S. Y. Lee, Y. J. Jo, Z. Wu, E. Y. Choi, R. J. Samulski, and P. A. Campochiaro. Overexpression of SOD in retina: Need for increase in H 2O 2-detoxifying enzyme in same cellular compartment. *Free Radic. Biol. Med.*, 51(7):1347–1354, 2011. [Cited on page 7.]
- S. Vakili, V. Picheny, and N. Durrande. Regret Bounds for Noise-Free Bayesian Optimization. *arXiv*, 2020. [Cited on page 115.]
- G. Valle, A. Mazzoni, F. Iberite, E. D'Anna, I. Strauss, G. Granata, M. Controzzi, F. Clemente, G. Rognini, C. Cipriani, T. Stieglitz, F. M. Petrini, P. M. Rossini, and S. Micera. Biomimetic Intraneural Sensory Feedback Enhances Sensation Naturalness, Tactile Sensitivity, and Manual Dexterity in a Bidirectional Prosthesis. *Neuron*, 100(1): 37–45.e7, 2018. [Cited on page 34.]
- J. J. van Rheede, C. Kennard, and S. L. Hicks. Simulating prosthetic vision: Optimizing the information content of a limited visual display. *J. Vis.*, 10(14):32, 2010. [Cited on page 33.]
- C. Veraart, C. Raftopoulos, J. T. Mortimer, J. Delbeke, D. Pins, G. Michaux, A. Vanlierde, S. Parrini, and M. C. Wanet-Defalque. Visual sensations produced by optic nerve stimulation using an implanted self-sizing spiral cuff electrode. *Brain Res.*, 813(1):181–186, 1998. [Cited on page 14.]
- C. Veraart, M. C. Wanet-Defalque, B. Gérard, A. Vanlierde, and J. Delbeke. Pattern Recognition with the Optic Nerve Visual Prosthesis. *Artif. Organs*, 27(11):996–1004, 2003. [Cited on pages 14 and 14.]
- M. Volpp, L. P. Fröhlich, K. Fischer, A. Doerr, S. Falkner, F. Hutter, and C. Daniel. Meta-Learning Acquisition Functions for Transfer Learning in Bayesian Optimization. *ICLR*, 2019. [Cited on page 147.]
- A. Volta. On the Electricity Excited by the Mere Contact of Conducting Substances of Different Kinds. In a Letter from Mr. Alexander Volta, F. R. S. Professor of Natural Philosophy in the University of Pavia, to the Rt. Hon. Sir Joseph Banks, Bart. K. B. P. R. S. *Philos. Trans. R. Soc. London Ser. I*, 90(jan):403–431, 1800. [Cited on page 13.]
- S. T. Walston, R. H. Chow, and J. D. Weiland. Direct measurement of bipolar cell responses to electrical stimulation in wholemount mouse retina. *J. Neural Eng.*, 15(4):046003, 2018. [Cited on pages 27 and 27.]
- J. Wang, X. Wu, Y. Lu, H. Wu, H. Kan, and X. Chai. Face recognition in simulated prosthetic vision: Face detectionbased image processing strategies. *J. Neural Eng.*, 11(4), 2014. [Cited on page 33.]
- L. Wang, K. Mathieson, T. I. Kamins, J. D. Loudin, L. Galambos, G. Goetz, A. Sher, Y. Mandel, P. Huie, D. Lavinsky, J. S. Harris, and D. V. Palanker. Photovoltaic retinal prosthesis: implant fabrication and performance. *J. Neural Eng.*, 9(4):046014, 2012. [Cited on page 24.]

- L. Wang, F. Sharifian, J. Napp, C. Nath, and S. Pollmann. Cross-task perceptual learning of object recognition in simulated retinal implant perception. *J. Vis.*, 18(13):22, 2018. [Cited on page 33.]
- Y. Wang, C. Wang, and W. Powell. The knowledge gradient for sequential decision making with stochastic binary feedbacks. *33rd Int. Conf. Mach. Learn. ICML* 2016, 3:1761–1773, 2016. [Cited on pages 94 and 121.]
- Z. Wang, C. Gehring, P. Kohli, and S. Jegelka. Batched Large-scale Bayesian Optimization in High-dimensional Spaces. Int. Conf. Artif. Intell. Stat. AISTATS 2018, 84:745–754, 2017a. [Cited on page 109.]
- Z. Wang, C. Li, S. Jegelka, and P. Kohli. Batched high-dimensional Bayesian optimization via structural kernel learning. 34th Int. Conf. Mach. Learn. ICML 2017, 8:5590–5603, 2017b. [Cited on pages 68 and 142.]
- H. Wässle. Parallel processing in the mammalian retina. Nat. Rev. Neurosci., 5(10):747–757, 2004. [Cited on page 5.]
- A. B. Watson. A formula for human retinal ganglion cell receptive field density as a function of visual field location.
 J. Vis., 14(7):1–17, 2014. [Cited on pages 2 and 28.]
- A. B. Watson. QUEST+: A general multidimensional Bayesian adaptive psychometric method. J. Vis., 17(3):10, 2017. [Cited on pages 56 and 62.]
- R. Weeks, B. Horwitz, A. Aziz-Sultan, B. Tian, C. M. Wessinger, L. G. Cohen, M. Hallett, and J. P. Rauschecker. A positron emission tomographic study of auditory localization in the congenitally blind. *J. Neurosci.*, 20(7): 2664–2672, 2000. [Cited on page 10.]
- D. Weichert and A. Kister. Bayesian Optimization for Min Max Optimization. *ICML 2020 Work. Real World Exp. Des. Act. Learn.*, pages 1–10, 2020. [Cited on page 131.]
- J. D. Weiland, S. T. Walston, and M. S. Humayun. Electrical Stimulation of the Retina to Produce Artificial Vision. *Annu. Rev. Vis. Sci.*, 2:273–294, 2016. [Cited on pages 26 and 27.]
- A. C. Weitz, D. Nanduri, M. R. Behrend, A. Gonzalez-Calle, R. J. Greenberg, M. S. Humayun, R. H. Chow, and J. D. Weiland. Improving the spatial resolution of epiretinal implants by increasing stimulus pulse duration. *Sci. Transl. Med.*, 7(318):1–12, 2015. [Cited on pages 29, 30, and 30.]
- C. G. Welker, A. S. Voloshina, V. L. Chiu, and S. H. Collins. Shortcomings of human-in-the-loop optimization of an ankle-foot prosthesis emulator: a case series. *R. Soc. Open Sci.*, 8(5):rsos.202020, 2021. [Cited on pages 41 and 72.]
- S. Whiteley, T. Litchfield, P. Coffey, and R. Lund. Improvement of the Pupillary Light Reflex of Royal College of Surgeons Rats Following RPE Cell Grafts. *Exp. Neurol.*, 140(1):100–104, 1996. [Cited on page 12.]
- T. N. Wiesel and D. H. Hubel. Spatial and chromatic interactions in the lateral geniculate body of the rhesus monkey. *J. Neurophysiol.*, 29(6):1115–1156, 1966. [Cited on page 15.]
- R. Wilke, V. P. Gabel, H. Sachs, K. U. B. Schmidt, F. Gekeler, D. Besch, P. Szurman, A. Stett, B. Wilhelm, T. Peters, A. Harscher, U. Greppmaier, S. Kibbel, H. Benav, A. Bruckmann, K. Stingl, A. Kusnyerik, and E. Zrenner. Spatial resolution and perception of patterns mediated by a subretinal 16-electrode array in patients blinded by hereditary retinal dystrophies. *Investig. Ophthalmol. Vis. Sci.*, 52(8):5995–6003, 2011. [Cited on pages 28 and 30.]
- S. J. Wilks, S. M. Richardson-Burns, J. L. Hendricks, D. C. Martin, and K. J. Otto. Poly(3,4-ethylenedioxythiophene) as a micro-neural interface material for electrostimulation. *Front. Neuroeng.*, 2(JUN):1–8, 2009. [Cited on page 25.]
- N. R. Williams and M. S. Okun. Deep brain stimulation (DBS) at the interface of neurology and psychiatry. *J. Clin. Invest.*, 123(11):4546–4556, 2013. [Cited on page 13.]
- A. G. Wilson, Z. Hu, R. Salakhutdinov, and E. P. Xing. Deep kernel learning. Proc. 19th Int. Conf. Artif. Intell. Stat. AISTATS 2016, 51:370–378, 2016. [Cited on page 143.]
- B. S. Wilson and M. F. Dorman. Cochlear implants: Current designs and future possibilities. J. Rehabil. Res. Dev., 45 (5):695–730, 2008. [Cited on page 13.]

- J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Efficiently Sampling Functions from Gaussian Process Posteriors. *arXiv*, 2020a. [Cited on pages 108, 109, 110, 112, 113, and 122.]
- J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Pathwise conditioning of gaussian processes. *arXiv*, 2020b. [Cited on page 109.]
- M. Wistuba and J. Grabocka. Few-Shot Bayesian Optimization with Deep Kernel Surrogates. In *ICLR*, pages 1–15, 2021. [Cited on page 146.]
- G. Woch, R. B. Aramant, M. J. Seiler, B. T. Sagdullaev, and M. A. McCall. Retinal transplants restore visually evoked responses in rats with photoreceptor degeneration. *Invest. Ophthalmol. Vis. Sci.*, 42(7):1669–76, 2001. [Cited on page 12.]
- W. L. Wong, X. Su, X. Li, C. M. G. Cheung, R. Klein, C. Y. Cheng, and T. Y. Wong. Global prevalence of age-related macular degeneration and disease burden projection for 2020 and 2040: A systematic review and meta-analysis. *Lancet Glob. Heal.*, 2(2):e106–e116, 2014. [Cited on page 8.]
- A. F. Wright, C. F. Chakarova, M. M. Abd El-Aziz, and S. S. Bhattacharya. Photoreceptor degeneration: Genetic and mechanistic dissection of a complex trait. *Nat. Rev. Genet.*, 11(4):273–284, 2010. [Cited on page 6.]
- J. Wu and P. I. Frazier. The parallel knowledge gradient method for batch Bayesian optimization. *Adv. Neural Inf. Process. Syst.*, 2016. [Cited on pages 73 and 153.]
- J. Wu and P. I. Frazier. Practical Two-Step Look-Ahead Bayesian Optimization. In *Neural Inf. Process. Syst.*, 2019. [Cited on pages 73, 107, 146, and 147.]
- K. Xue and R. E. MacLaren. Ocular gene therapy for choroideremia: clinical trials and future perspectives. *Expert Rev. Ophthalmol.*, 13(3):129–138, 2018. [Cited on page 12.]
- D. L. K. Yamins and J. J. DiCarlo. Using goal-driven deep learning models to understand sensory cortex. Nat. Neurosci., 19(3):356–365, 2016. [Cited on page 38.]
- G. N. Yannakakis and J. Hallam. Ranking vs. preference: A comparative study of self-reporting. Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics), 6974 LNCS(PART 1):437–446, 2011. [Cited on page 42.]
- J. C. Young and C. E. Minder. Algorithm AS 76: An Integral Useful in Calculating Non-Central t and Bivariate Normal Probabilities. *Appl. Stat.*, 23(3):455, 1974. [Cited on pages 53 and 88.]
- R. W. Young and D. Bok. Participation of the retinal pigment epithelium in the rod outer segment renewal process. J. Cell Biol., 42(2):392–403, 1969. [Cited on page 5.]
- D. Y. Yu, S. Cringle, K. Valter, N. Walsh, D. Lee, and J. Stone. Photoreceptor death, trophic factor expression, retinal oxygen status, and photoreceptor function in the P23H rat. *Investig. Ophthalmol. Vis. Sci.*, 45(6):2013–2019, 2004. [Cited on page 6.]
- C. Zeitz, A. G. Robson, and I. Audo. Congenital stationary night blindness: An analysis and update of genotypephenotype correlations and pathogenic mechanisms. *Prog. Retin. Eye Res.*, 45:58–110, 2015. [Cited on page 31.]
- J. Zhang, P. Fiers, K. A. Witte, R. W. Jackson, K. L. Poggensee, C. G. Atkeson, and S. H. Collins. Human-in-theloop optimization of exoskeleton assistance during walking. *Science* (80-.)., 356(6344):1280–1284, 2017. [Cited on page 40.]
- M. Zhang, H. Li, and S. Su. High Dimensional Bayesian Optimization via Supervised Dimension Reduction. *arXiv*, pages 4292–4298, 2019. [Cited on pages 68 and 142.]
- S. Zhao, Y. Feng, T. Wu, and C. Yang. Application of Biocompatible Organic/Inorganic Composite Film in Implantable Retinal Prosthesis. In 2018 IEEE Int. Conf. Cyborg Bionic Syst., pages 101–105. IEEE, 2018. [Cited on page 26.]

- M. Zoghi, Z. Karnin, S. Whiteson, and M. De Rijke. Copeland dueling bandits. In *Adv. Neural Inf. Process. Syst.*, volume 2015-Janua, pages 307–315, 2015. [Cited on page 90.]
- E. Zrenner. Will retinal implants restore vision? Science (80-.)., 295(5557):1022-1025, 2002. [Cited on page 16.]
- E. Zrenner. Fighting blindness with microelectronics. Sci. Transl. Med., 5(210), 2013. [Cited on page 14.]
- E. Zrenner, R. Wilke, T. Zabel, H. Sachs, K. Bartz-Schmidt, F. Gekeler, B. Wilhelm, U. Greppmaier, A. Stett, and S. S. Group. Psychometric analysis of visual sensations mediated by subretinal microelectrode arrays implanted into blind retinitis pigmentosa patients. *Invest. Ophthalmol. Vis. Sci.*, 48(13):659, 2007. [Cited on page 13.]
- E. Zrenner, K. U. Bartz-Schmidt, H. Benav, D. Besch, A. Bruckmann, V. P. Gabel, F. Gekeler, U. Greppmaier, A. Harscher, S. Kibbel, J. Koch, A. Kusnyerik, T. Peters, K. Stingl, H. Sachs, A. Stett, P. Szurma, B. Wilhelm, and R. Wilke. Subretinal electronic chips allow blind patients to read letters and combine them to words. *Proc. R. Soc. B Biol. Sci.*, 278(1711):1489–1497, 2011. [Cited on pages 16, 21, 23, and 26.]