Development and validation of stimulation strategies for the optogenetics

Quentin Sabatier

To cite this version:

Quentin Sabatier. Development and validation of stimulation strategies for the optogenetics. Automatic. Sorbonne Université, 2018. English. NNT : 2018SORUS083 . tel-02498147

HAL Id: tel-02498147
https://theses.hal.science/tel-02498147
Submitted on 4 Mar 2020

Sorbonne Université

Ecole doctorale Sciences mécaniques, acoustique, électronique & robotique de Paris
Institut de la Vision / Equipe Vision et Calcul Naturel

Conception and Validation of Stimulation Strategies for the Optogenetics

Thèse CIFRE en collaboration avec GenSight Biologics

Par Quentin Sabatier

Thèse de doctorat de Robotique

Dirigée par Pr. Ryad Benosman

Présentée et soutenue publiquement le 7 mai 2018

Devant un jury composé de :

Dr.	Yves Frégnac	Rapporteur
Dr.	Laurent Perrinet	Rapporteur
Prof.	Serge Picaud	Examinateur
Prof.	Stéphane Régnier	Examinateur
Associate Prof.	Sio-Hoï Ieng	Examinateur
Prof.	Ryad Benosman	Directeur de thèse

Quentin Sabatier: *Conception and Validation of Stimulation Strategies for the Optogenetics,* Thèse CIFRE en collaboration avec GenSight Biologics, April 16, 2018

SUPERVISORS:

Pr. Ryad Benosman — PhD Advisor
Didier Pruneau, PhD — Industrial supervisor

LOCATION:

Institut de la Vision
17, rue Moreau, 75012 Paris
GenSight Biologics
74, rue du Faubourg Saint-Antoine, 75012 Paris

ABSTRACT

1.5 million people throughout the world suffer from Retinitis Pigmentosa, a group of inherited diseases resulting in degeneration of the retina. In turn, this leads to loss of night vision, progressive narrowing of the visual field, and eventually total blindness. Due to the heterogeneity in the mutations causing the disease, solutions to compensate for the symptoms of the disease are being investigated. These systems, called *retinal prostheses*, usually consist in three main elements: (i) a camera filming the scene in front of the patient, usually mounted on a pair of goggles worn by the patient and (ii) a stimulation device that is able to elicit activity in the patient's nervous system (often at the level of the retina) and (iii) a processor unit that monitors the transformation between the signal output by the camera and the stimulation patterns.

The work presented in this thesis is part of the effort of GenSight Biologics to develop such retinal prosthesis. It combines two emerging technologies, a *neuromorphic* camera in which each pixel acquires the time-varying signal in an asynchronous data-driven fashion and *Optogenetics* which consists in rendering targeted neurons photoexcitable.

My work spans across the whole processing pipeline. In Chapter 2, we present a algorithm extracting the time-varying spatial frequencies based on the asynchronous stream of events output by the camera. This algorithm aims at simplifying the scene before it is shown to the patient. This step is required because limited information can flow through the Brain-Computer Interface (the point of contact between the hardware and the nervous system of the patient) and patients equipped with retinal prostheses have difficulties coping with the richness of visual scenes found in the real world, even if they perform well in controlled tasks within the lab setting.

Chapter 3 and Chapter 4 focus on the Brain-Computer Interface itself by proposing a model of the transformation between the light signal projected from the goggles and the trains of Action Potentials triggered by the Retinal Ganglion Cells of the patient and carrying the information to downstream areas of the visual pathway. Chapter 3 models the response of ChrimsonR — the protein underlying the photoexcitability of the Retinal Ganglion Cells — to the light signal, while Chapter 4 describes the process leading to the generation of Action Potentials. Modeling the Interface has the final purpose to use its capacities to its full potential, and maximize the rate of information from the visual scene that can be communicated to the patient.

RÉSUMÉ

A travers le monde, 1.5 millions de personnes souffrent de la Rétino-pathie Pigmentaire, une famille de maladies héréditaires entraînant une dégénérescence de la rétine. Les premiers symptômes sont la perte de la vision nocturne et du champ visuel périphérique. Par la suite, le rétrécissement du champ visuel s'accentue et mène à une cécité totale. En raison de l'hétérogénéité des mutations génétiques responsables de la maladie, des solutions visant à compenser les symptômes de la maladie émergent. Parmi ces systèmes, les *prothèses rétiniennes* comportent trois éléments principaux: (i) une caméra filmant la scène devant le patient, elle est habituellement montée sur une paire de lunettes et (ii) un dispositif de stimulation qui est capable de contrôler une partie de l'activité neuronale du patient (souvent au niveau de la rétine) et (iii) un processeur qui implémente la transformation entre le signal de sortie de la caméra et les commandes de stimulation.

Le travail présenté dans cette thèse contribue au travail de GenSight Biologics pour développer une telle prothèse rétinienne. Le project combine deux technologies récentes, une caméra *neuromorphique* dans laquelle chaque pixel acquiert le signal d'une manière asynchrone, et une très haute résolution temporelle, et l'*optogénétique* qui permet de rendre les neurones ciblés photoexcitables.

Mon travail s'étend sur l'ensemble de la chaîne de traitement du signal. Dans le chapitre 2, nous présentons un algorithme extrayant les fréquences spatiales de la vidéo à partir du flux de mesures asynchrones émises par la caméra. Cet algorithme sert à simplifier la scène visuelle avant qu'elle ne soit communiquée au patient. Cette étape est nécessaire car la quantité d'information pouvant être transmise au travers de l'Interface Cerveau-Machine est limitée; et les patients équipés de prothèses rétiniennes ont souvent des difficultés à faire face à la richesse des scènes visuelles réelles, même s'ils réussissent bien les tâches auxquelles ils sont confrontés dans l'environnement contrôlé du laboratoire.

Les chapitres 3 et 4 se concentrent sur l'Interface Cerveau-Machine en développant un modèle de la transformation reliant le signal lumineux projeté par les lunettes et les trains de potentiels d'action déclenchés par les cellules ganglionnaires de la rétine du patient. Ces potentiels d'action sont le support de l'information communiquées vers les zones responsables de la perception visuelle. Le chapitre 3 modélise la réponse de ChrimsonR — la protéine qui sous-tend la photoexcitabilité induites dans les cellules ganglionnaires de la rétine — à la commande lumineuse, tandis que le chapitre 4 décrit le processus de déclenchement des potentiels d'action. La modélisation de l'Interface Cerveau-Machine a pour objectif final d'utiliser ses capacités à son plein potentiel et de maximiser le taux d'information à propos de la scène visuelle qui peut être communiqué au patient.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACRONYMS & NOTATIONS

RP	Retinitis Pigmentosa
AMD	Age-related Macular Degeneration
BCI	Brain-Computer Interface
DNA	Deoxyribonucleic acid
AAV	Adeno-associated Virus
IML	Inner Limiting Membrane
DMD	digital micro-mirror device
DVS	Dynamic Vision Sensor
ATIS	Asynchronous Time-based Image Sensor
CNS	Central Nervous System
AP	Action Potential
RPE	Retinal Pigment Epithelium
IPL	Inner-Plexiform Layer
OPL	Outer-Plexiform Layer
RGC	Retinal Ganglion Cell
DSGC	Direction-Selective Ganglion Cell
LGN	Lateral Geniculate Nucleus
HEK	Human Embryonic Kidney
ChR-1	channelrhodopsin-1
ChR-2	channelrhodopsin-2
L-LN-IF	Linear-NonLinear Integrate-and-Fire
SDF	Spike-Density Function
PSTH	Peri-Stimulus Time Histogram

Part I

INTRODUCTION

1

INTRODUCTION

1.1 RETINITIS PIGMENTOSA

Retinitis Pigmentosa (RP) is a group of inherited diseases that affect 1.5 million people worldwide and approximately 400.000 in the U.S. and in Europe, with a prevalence of 1 in 4000[60]. This disease leads to progressive vision loss through death of the rod photoreceptors, followed by the death of the cone photoreceptors.

There is a great number of genes responsible for the occurrence of RP, with a variety of inheritance patterns : autosomal-dominant (about 30-40% of cases), autosomal-recessive (50-60%) or X-linked (5-15%). More complex, non-mendelian, inheritance patterns have also been reported[75] (cited by Hartong et al.). To date, more than 45 genes causing the disease have been identified, and they only account for 60% of the patient population[60]. Most of the known genes affect the metabolism of rods or retinal pigment epithelium[124], explaining why rods die first. It seems that the cones are not directly affected by the gene deficiency, but rather die because of the absence of the rods. Being able to have the cones survive would already be a great success in tackling RP. However, the mechanism leading to cones death is not precisely identified[98]. To date, nutritional complements such as vitamin A are prescribed in order to slow down the loss of cone photoreceptors.

Usually, the symptoms start by a decrease in night vision and loss of far peripheral vision during childhood. These first signs can be difficult to detect by the patients, and it is difficult to assess whether the variability in detecting the disease (it can range from early childhood to adulthood[60]) is caused by a variability in the actual onset of the disease or in the detection by the patient. The disease progressively narrows the visual field, leading to *tunnel vision*. *Tunnel vision* is characterized by a very small field of view, and corresponds to full loss of the rod photoreceptors and survival of the central region of the retina, called the *fovea*, in which there are no rods and the highest density of cones. During this period, patients can maintain a high visual acuity in this central region. Eventually, all cones die out leading to complete blindness towards the age of 60[60].

There is no cure for RP and current therapies mostly aim to slow down cell death and vision loss. Nutritional supplements such as vitamin A have been used to prevent disease progression[22] as well as intravitreal injection of growth and/or survival factors[98]. However these treatments are very limited in blocking or reversing disease progression.

Gene therapy is under investigation as reviewed by Sahel and Roska. It can take one of two forms : the first one — gene supplementation or gene replacement — depends on the mutated gene or on the mutation itself, whereas the second one only depends on the structural and functional state of the retina. The first approach is the more desirable as it actually cures the disease. However, this would require developing a treatment for each gene responsible for every single form of the disease. Mutation-independent strategies are thus more promising. They include *neuroprotection* (attempting to slow down retinal degeneration) and strategies such as (i) electronic implant[66], (ii) stem cell therapy[129, 106] and (iii) *optogenetic* approaches[28].

If successful, mutation-independent approaches could be translated to patients suffering from other similar diseases, such as Age-related Macular Degeneration (AMD), a disease affecting approximately 8 million people worldwide[44].

1.2 STRUCTURE OF THE VISION RESTORATION DEVICE

The work presented here is part of the effort of the company GenSight Biologics — in collaboration with which I was doing my PhD research — to develop retinal prostheses compensating for the loss of vision in RP patients. This project is called GS030, and this is how we will be referring to it from now on. The rationale behind the project is that, although the photoreceptors have died, the four other layers of the retina are still alive and functional. The objective is to replace the photoreceptor layer with an electronic device standing in for the photoreceptors.

From an information processing point of view, photoreceptors have a double role: (i) they acquire the light signal, i.e. they modulate their own electrochemical behavior according to the time-varying properties (mostly intensity and spectral content) of the incoming light and (ii) they transmit this information to downstream layers of the retina, i.e. their electrochemical behavior drives the electrochemical behavior of the neurons they synapse with. Consequently, building a retinal prosthesis consists in developing a device that performs these two functions.

In order to do so, retinal prostheses usually consist in the same set of elements: (i) a sensor acquiring the visual information, (ii) a Brain-Computer Interface (BCI) stimulating the nervous system of the patient according to the sensor's output and (iii) a processing unit responsible for the mapping between the camera's output and the stimulation commands sent to the BCI.

In GS030, the choice for the sensor is the Asynchronous Time-based Image Sensor (ATIS), a camera inspired by the retina which is sampling the visual scene in a non-synchronized fashion. Its design results in the acquisition of visual scenes with a very high temporal resolution (of the order of the millisecond), low data load and high dynamic range. Its development is part of the wider field of neuromorphic engineering, a line of research aiming at building hardware and software closer to the way the brain processes information than current computers do.

The Brain-Computer Interface is implemented through a technology called *Optogenetics*. It consists in expressing light-sensitive proteins in naturally non photosensitive neurons in order to make them photoexcitable. The light-sensitive proteins are naturally expressed in micro-organisms, and subtend their vision. Thanks to years of research, the gene sequences responsible for the expression of several of such proteins are known and can be used to have these proteins expressed in other organisms and other cells. In GS030, the gene sequence is going to be carried by a viral capsid, an Adeno-Associated Virus injected into the vitreous humor of the eye of the patient, up to the retina where it is expected to transfect mainly the Retinal Ganglion Cells, the output layer of the retina.

The targeted Retinal Ganglion Cells will then become photoexcitable, and a light projector mounted on goggles worn by the patient will allow direct control of the electrochemical behavior of these cells. Most importantly, focusing light on the ganglion cells will induce action potentials, fast electrical impulses that subtend information communication in the brain. These action potentials will travel along the axons of the ganglion cells, forming the optic nerve, the line of communication from the retina to the rest of the visual pathway.

The ATIS camera is fixed on the goggles, and therefore acquires the visual scene right in front of the patient. The projector consists in a light source and a digital micro-mirror device (DMD), a 240 × 320 array of micro-mirrors which have two positions, an ON position reflecting the light towards the retina and an OFF position sending the light towards an absorbent material. Finally, a processor unit allows some flexibility in the way the position of the micro-mirrors are controlled according to the sensor's output signal[55].

Within the whole project, the work presented here focuses on information processing and aims at designing algorithms running on the processor unit carrying out the transformation between the sensor and the light projector with the final intention that the information encoded in the spike trains of the Retinal Ganglion Cells of the patient reflect the relevant information from the visual scene.

1.3 SENSORY SUBSTITUTION

The problem of restoring vision in congenitally blind and late-blind people has been tackled as soon as the 60's through a method named

Sensory Substitution, with the seminal work of Paul Bach-y-Rita (1934 - 2006)[12]. Sensory substitution consists in *restoring* a sense by providing the information related to this sense through another sensory pathway. A successful example consisted in providing visual information through electrochemical stimulation on the tongue[13].

The major difference between sensory *substitution* and *restoration* is the fact that sensory restoration uses the original sensory pathway to transmit information. However, due to the limitations at the interface between the stimulation device and the nervous system, the field of sensory restoration faces a great number of challenges which are shared with sensory substitution. Sensory restoration can thus benefits from the lessons and experience gathered during decades of research on sensory substitution.

In particular, sensory substitution systems provided major insight into brain plasticity. One of the most striking results was that sensory modalities that are still functional could recruit neurons in lost sensory pathways (e.g. [116]). Additionally, a few people using sensory substitution devices extensively have been able to take advantage of the device to perform some every-day life tasks, as well as advanced visual tasks such as depth perception[140]. Overall, the results obtained in sensory substitution have surpassed the expectations.

However, as pointed out by Maidenbaum et al. and Elli et al., the results have mostly been obtained in lab settings as part of experiments. The vast majority of subjects dit not manage to benefit from their device to perform everyday life tasks, even if they performed well on the tasks which were presented during the study.

In the remaining of this section, we list a number of teachings from sensory substitution which are of direct interest for our optogenetics-based vision restoration project and for my work in particular.

A major conclusion is that the subject must control the sensor acquiring the data, e.g. the camera in a context of sight restoration, so that the learning phase of the new sense can occur[10]. Through this control, the patient can understand how its actions on the sensor modify his perception and thus understand how the new perception works. In the case of vision restoration, Bach-y Rita suggest that the patient has control over the position, the aperture and focus of the camera. In the GS030 project, the camera is mounted on goggles worn by the subject and therefore the subject controls the position of the camera through head motion.

The field of sensory substitution also established the general structure of information processing in the device. The pipeline is divided in three blocks: the sensor, a preprocessing step and the BCI. According to Bach-y Rita, the major challenge lies in the BCI. Optogenetics offers the possibility (in a near future) to provide an interface with a spatial resolution of the order of the single cell and a temporal resolution ranging from 1 ms to 10 ms. The efficiency of the interface relates to the issue of *sensory overload*, the fact that the subject cannot cope with too much information if not transmitted properly[11]. Eventually, the patient does not perceive enough information, which deteriorates the efficiency of the strategy as a whole. Consequently, optimizing the transmission rate at the level of the BCI is a major issue in sensory substitution and restoration projects, and is the focus of Chapter 3 and Chapter 4.

The preprocessing step is another major issue. In fact, it is a great challenge for sensory substitution devices to move from the lab to the *real* world, which usually constitutes a much richer sensory environment[43]. In this situation, and since the GS030 device might face the same issue, it would be interesting to develop a preprocessing step that does not only map the acquired information to the range of perceivable stimulation, but to consider extracting relevant information from the visual scene in order to limit the quantity of information transmitted through the interface. In Chapter 2, we developed an algorithm extracting spatial frequency information from the visual scene, which constitutes a possible choice for the preprocessing step.

Finally, and even if it does not impact our work directly, much effort has been devoted to developing rehabilitation programs, i.e. the familiarization process with the device. This process is decisive for the cerebral plasticity mechanism to take place in the subject's brain. This step is paramount. It is also very demanding for the subject, which suggests that as much work as possible should be done before going to the clinical trials and avoid using patient time to fix problems that could be solved through other means.

Compared to existing substitution devices, the strategy set up in GS030 seems to offer several advantages. The first one, shared with retinal implant based strategies, is to stimulate the natural visual pathways. We hope this configuration results in two beneficial effects. First, these circuits are designed to cope with the large amount of data associated with vision and thus very well suited to cope with the new rich stream of information. Second, visual information will travel along the same pathways as in natural vision and therefore allow the downstream structures involved in natural vision processing to perform their job in the best possible conditions.

1.4 BUILDING VISUAL PROSTHESES

In the GS030 project, information processing is implemented using the same structure as the one developed for sensory substitution: a sensor, a preprocessing step and the BCI. The sensor is a bio-inspired asynchronous camera. The camera is presented in Section 1.5 and acquires the visual scene with a millisecond temporal resolution. In order to achieve this while generating a sensible amount of data, the output of the camera consists in a stream of measures which are not synchronized across pixels. This format for the data — and in order to maintain the high-temporal resolution throughout the whole processing pipeline up to the BCI — imposes a certain structre, termed event-driven. This algorithmic paradigm is widely used in the neuromorphic engineering field, presented in Section 1.6. Finally, the Brain-Computer Interface is implemented using the technic called *Optogenetics*. It consists in genetically expressing light-sensitive transmembrane proteins, called *channelrhodopsins*, in neurons. The targeted neurons are then photo-excitable, which means that their electro-chemical activity (the substrate of information processing in the nervous system) can be driven by external light. This technic has been widely used in the neuroscience community, since it has been first expressed in neuron in 2005, and is presented in Section 1.7.

1.5 THE ASYNCHRONOUS TIME-BASED IMAGE SENSOR

Development of the ATIS, the sensor chosen in the GS030 project, is part of the field called neuromorphic engineering. Since the development of the first *Formal Neuron* model in 1943 by Warren McCulloch and Walter Pitts, who were studying the computational analogies between brains and computers, an important effort has been devoted to creating computers, and sensors, which work more like the brain. In fact, it quickly became clear that the nervous systems — human, but not only — were surpassing computers in a vast number of tasks[36].

With this in mind, Carver Mead developed a *Silicon Retina*[97], a camera mimicking certain aspects of the function of the mammalian retina. Overcoming a number of difficulties, the Dynamic Vision Sensor (DVS) was the first practical neuromorphic image sensor[87]. Its pixels function as independent temporal contrast detectors. Each one of them tracks the light intensity at its level, and remains silent while it stays in a close interval around the *current* value. Whenever the intensity crosses either threshold of the interval, an *event* is triggered (with an ON polarity if the upper threshold has been crossed, i.e. light intensity increases, and an OFF polarity otherwise) and output by the camera. Events are an important notion in the field of neuromorphic engineering, and are very close to the action potentials triggered by biological neurons and underlying information communication in the nervous system.

The ATIS shares the temporal contrast pixels with the DVS, but they are coupled with another kind of pixels performing an absolute measure of light intensity. In fact, temporal contrast pixels never provide any information relative to the absolute light level, they only provide information about changes. In the ATIS, absolute measurement pixels are triggered by the temporal contrast pixels. Absolute measures are only performed, and communicated, if relevant, i.e. if light intensity at the location of the pixel has changed since the last measure.

The choice of the ATIS as part of the vision restoration device follows the recent observation that temporal resolution is an important feature of vision processing in the nervous system. In mammal early vision processing structures, such as the retina[21, 138, 109, 135], the Lateral Geniculate Nucleus (LGN)[138, 120, 90, 82] or the Middle Temporal (MT) area of the visual cortex[27], action potentials triggered in response to a visual stimulus have a temporal precision close to 1 ms. Moreover, Butts et al. showed that this precision is driven by the stimulus, suggesting information is encoded in a *temporal code*[30, 29]. Since then, it has been shown human observers can make use of high temporal resolution displays (1 kHz equivalent frame rate) when performing discrimination and detection tasks[77].

This suggests that a significant amount of visual information lies in the fine temporal details of the signal. The ATIS is able to acquire the visual scene with this fine temporal scale, with low power consumption and low data generation, which are an absolute requirement for any embedded device such as the GS030 vision restoration device. The hope is that the millisecond resolution can be maintain throughout the whole processing pipeline and through the BCI, knowing that the visual system naturally copes with this type of resolution and therefore that the patient's nervous system will be able to benefit from this additional information, without creating any sensory overload.

1.6 PRE-PROCESSING: EVENT-BASED COMPUTING

Owing to the ATIS data format, parallel asynchronous streams of temporal contrast events and absolute measurements, the processing steps cannot be carried out using conventional computer vision algorithms, which operate on *frames*, i.e. the synchronized brightness measures for all pixels. Instead, algorithms must be organized in formal neuron networks. In these networks, formal neurons are elementary computational units which perform simple operations whenever they receive an event.

This computational paradigm is *distributed* and *event-driven*. Distributed means that, instead of having a central agent orchestrating all operations, each computational unit performs its own operations independently, based on the data it receives and a set of rules. Event-driven refers to the fact that neurons only react to incoming events. When receiving an event, either from another neuron or directly from the camera, the neuron performs an operation possibly on the information carried by the event (e.g. the polarity of the temporal contrast event) and on internal information. This operation can result in triggering one or several events which are then sent to connected neurons.

Event-driven computing yields the best results when implemented on adapted hardware. In fact, efficient event-driven computing requires the computation to be performed independently in parallel, and much effort is currently devoted by the neuromorphic community and by the industry to develop highly parallel, neuromorphic-oriented, hardware[68, 54, 53].

However, designing efficient event-driven algorithms is not such simple task, even when powerful adapted hardware is available. Having the computations to be distributed and event-driven are strong constraints. Algorithmic development therefore constitutes a significant part of the work in neuromorphic engineering[17, 67, 85, 41, 32, 142].

1.7 THE BRAIN-COMPUTER INTERFACE

As introduced in Section 1.3, the BCI is a crucial aspect of the project. In the GS030 project, the interface is implemented using *Optogenetics*. This technology relies on gene therapy, the therapeutic delivery of Deoxyribonucleic acid (DNA) to a patient. The delivery can be achieved through multiple methods, including transport by a recombinant virus and injection of naked DNA. The final goal is to have the transfected cell express proteins based on the new genetic material as if it was its own.

In this context, the following section presents the field of optogenetics and the one after focuses on the protein which is being expressed as a result of expression of the transfected genetic sequence. Finally, *Targeting the eye* describes the specific transfection procedure which has been chosen as part of the GS030 project.

Optogenetics

Optogenetics originates from the study of vision in microalgae[62]. It led to the discovery and study of the protein underlying vision in these organisms. The associated gene sequences of this family of proteins, such as channelrhodopsin-1 (ChR-1) and channelrhodopsin-2 (ChR-2), were then unveiled and the proteins were successfully expressed in other organisms[101, 102].

In 2005, Boyden et al. reported the successful expression of ChR-2 in mammalian neurons, and the ability to control the spiking activity of the ChR-2-expressing neurons with a temporal resolution of the order to the millisecond[24].

Since then, a number of practical details have been addressed and the technology has been widely adopted by the neuroscience community as a major tool[23, 40, 5, 7]. In parallel, channelrhodopsins have been studied intensively: their crystallographic structure has been established[74], mutant versions have been developed[121] and studies of their functional behavior is under intense investigation using spectroscopic[94] and simulation[18, 105, 58] methods.

ChrimsonR

ChR-2 is the most used and most studied channelrhodopsin. However, its peak spectral sensitivity is located in the blue at approximately 450 nm[128]. This does not make it a good candidate for clinical application because of the increased toxicity of this portion of the light spectrum. Efforts to shift the sensitivity spectra of channelrhodopsins through mutagenesis has proven difficult.

Instead, Klapoetke et al. looked for new wild type channelrhodopsins by sequencing the transcriptome of 127 species of alga. They discovered a large amount of new interesting proteins, including *Chronos* which proved to exhibit both the highest dynamics and high level of activation and *Chrimson* which is 45 nm more red-shifted than any other previously known channelrhodopsin, mutants included.

ChrimsonR — the chosen protein in GS030 — is an engineered mutant of Chrimson with similar spectral sensitivity and enhanced dynamics[78]. The study showed that the spiking activity of ChrimsonR-expressing neurons could be driven reliably through light excitation at 625 nm.

Targeting the eye

In [124], Sahel and Roska — both among the founders of GenSight Biologics — discuss the use of gene therapy to address retinal degeneration diseases and symptoms. Their conclusion is that "The eye has unique characteristics, compared with other tissues and organs, that make it particularly suited for gene therapy". The detailed reasons are :

- The eye is a small and closed compartment. Consequently, high viral concentration over a long time interval can be achieved with relatively small amount of injected virus.
- The different structures within the eye are separated from each other by protective membranes. Specific ocular structure can therefore be targeted by injecting the virus in the correct compartment. However, this advantage should be tempered in our case, since the target is the retina while the virus is injected in the vitreous humor, both spaces being separated by the Inner Limiting Membrane (IML). This aspect is discussed in more details below.
- Most cell types in the eye are stable and conserved across mammals. This suggests that preliminary studies on animal models are good indicators of success in humans. Moreover, most cell types

do not divide, ensuring that the delivered transgene will not be diluted into cell division.

- The eye is partially protected from the actions of the immune system by the *blood-retinal barrier*. This, and additional mechanisms, yield what is known as the *ocular immune privilege*. This ensures a large inhibition of the immune response directed against the gene products or the vector antigens.
- The transparency of the eye allows direct non-invasive assessment of the transfection process. It is obviously a huge asset when the final goal is light excitation of the transfected cells.
- In human clinical trials, the untreated eye can serve as control to assess the efficiency of the therapy in the treated eye.

When targeting the retina, the Adeno-associated Virus (AAV) can either be injected into the subretinal space (i.e. between the retina and the Retinal Pigment Epithelium (RPE) to which the retina is attached in natural conditions) or in the intravitreal space. In GS030, the intravitreal injection route has been selected. Even if not being the same ocular compartment as the retina, this strategy has several benefits, including (i) the vitreous humor is easily accessible, making the injection simple and (ii) it prevents the risk of retinal detachment associated with subretinal injection.

Because of the IML separating the retina from the vitreous humor, targeting the retina through intravitreal injection required the development of vectors specifically design for this purpose. Retinal Ganglion Cell (RGC) in the perifoveal regions are the primary targets of this kind of procedure, because the IML is weaker in the perifoveal region and that the RGC constitute the first nerve cell layer coming from the intravitreal cavity. Nevertheless, Dalkara et al. engineered AAV2.7m8, an AAV variant, designed to achieve higher transfection efficacy through all five layers of the retina and across the whole retina. In [39], they reported such successful transfection pattern in rodent and nonhuman primate retinas.

As part of GS030, four different construct were tested: two versions of the transgene, ChrimsonR alone and a combination ChrimsonR-tdTomato (a fluorescent protein) transported by two different viral capsids, either the wild type AAV2 or the variant AAV2.7m8. Using both imaging methods and electrophysiological recordings, the ChrimsonR-tdTomato construct combined with the AAV2.7m8 capsid showed to yield the best results, and was therefore chosen.

Immunohistochemical detection of tdTomato showed that ChrimsonR-tdTomato was essentially expressed in the Retinal Ganglion Cells located in the perifoveal ring. Epifluorescence and two-photon imaging revealed very bright soma with occasional dendritic arborisation. The observed cell morphologies suggest that most of the transfected cells are Midget Ganglion cells. This fact is not surprising since this cell type constitutes up to 95% of the RGC population in this area[37].

Part II

PREPROCESSING STAGE

2

ASYNCHRONOUS EVENT-BASED FOURIER TRANSFORM

2.1 INTRODUCTION

Conventional imaging devices sample scenes at a fixed frequency; all pixels acquire luminance simultaneously by integrating the amount of light over a fixed period of time. Often only very few pixels change between two consecutive frames, leading to the acquisition of large amounts of redundant data. Often only very few pixels change between two consecutive frames, leading to the acquisition of large amounts of redundant data. When a conventional frame-based camera observes a dynamic scene, no matter where the frame rate is set to, it will always be wrong because there is no relation whatsoever between dynamics present in a scene and the chosen frame rate, over-sampling and/or under-sampling occur, and moreover both usually happen at the same time. When acquiring a natural scene with a fast moving object in front of static background with a standard video camera, motion blurring and displacement of the moving object between adjacent frames will result from under-sampling the fast motion, while repeatedly sampling and acquiring static background over and over again. This will lead to large amounts of redundant, previously known data that do not contain any new information. As a result, the scene is simultaneously under- and over-sampled. This strategy of acquiring dynamic visual information has been accepted by the machine vision community for decades, likely due to the lack of convincing alternative.

An alternative to fixed-frequency is to sample a time-varying signal not on the time axis but using its the amplitude axis, leading to non uniform sampling rates that match the dynamics of the input signal. This sampling approach is often referred to as asynchronous delta modulation [61] or continuous-time level-crossing sampling [137]. Recently, this sampling paradigm has advanced from the recording of 1-D signals to the real-time acquisition of 2-D image data. The asynchronous time-based image sensor (ATIS) described in [115] contains an array of autonomously operating pixels that combine an asynchronous level-crossing detector and an exposure measurement circuit. Each exposure measurement by an individual pixel is triggered by a level-crossing event measuring a relative illuminance change at the pixel level. Hence, each pixel independently samples its illuminance, through an integrative measurement, upon detection of a change of a certain magnitude in this same illuminance, establishing its instantaneous gray level after it has changed. The result of the exposure measurement (i.e., the new gray level) is asynchronously transmitted off the sensor together with the pixels xy-address. As a result, image information is not acquired frame-wise but conditionally only from parts in the scene where there is new information. Only information that is relevant — because unknown — is acquired, transmitted and processed.

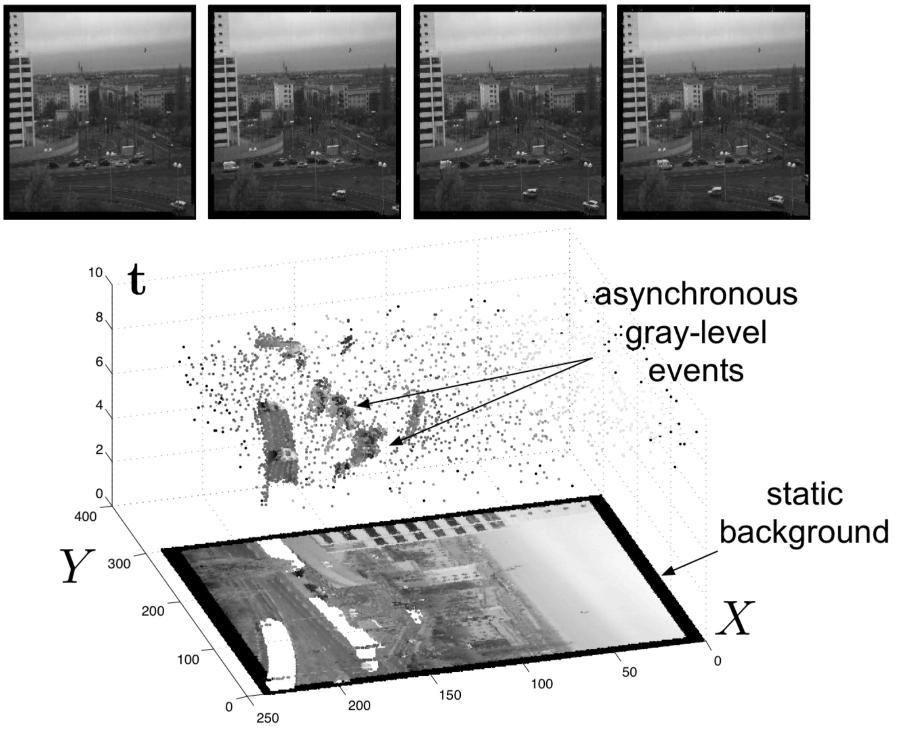

Figure 1: **The spatio-temporal space of imaging events**: Static objects and scene background are acquired first. Then, dynamic objects trigger pixel-individual, asynchronous gray level events after each change. Frames are absent from this acquisition process. Samples of generated images from the presented spatio-temporal space are shown in the upper part of the figure.

Fig. 1 shows the general principle of asynchronous imaging spaces. Frames are absent from this acquisition process. They can however be reconstructed, when needed (e.g. for display purposes), as shown at the top part of Fig. 1 and at frequencies limited only by the temporal resolution of the pixel circuits (up to hundreds of kiloframes per second). Static objects and background information, if required, can be recorded as a snapshot at the start of an acquisition henceforward moving objects in the visual scene describe a spatio-temporal surface at very high temporal resolution shown in the bottom part of Fig. 1. This novel paradigm of visual data acquisition calls for a new methodology in order to efficiently process the sparse, event-based image information without sacrificing its beneficial characteristics. Several methods have been recently published, which outperform conventional approaches both in computational costs and robustness. These cover all topics of machine vision: stereovision[79, 16, 122, 48], object recognition[107, 110], optical flow[17, 123], robotics[35, 103, 20, 25], tracking[104, 84] image processing[67] and retina prosthetics[92, 91].

This paper contributes to the field of asynchronous event-based vision by proposing an algorithm that computes the spatial Discrete Fourier Transform (DFT) iteratively for each incoming high temporal resolution event (1 µs time precision). The method computes the exact spatial DFT on event-based visual signals, without the need to reprocess already acquired information. This work also extends the event-based formulation of the DFT by introducing a lossy transformation methodology that can reduce even more computations by estimating a trade off between the quality of a reconstructed signal and the processing time. The time-varying visual signals are provided by the Asynchronous Time-based Image Sensor (ATIS) [115]. Conventional frame-based algorithms cannot be applied unchanged to event-based representation without leading to an immediate loss of the benefits inherent to the new sensing paradigm. Discrete Fourier Transform (DFT) is widely used in digital signal processing and scientific computing applications. The two-dimensional (2D) DFT is used in a wide variety of imaging applications that need spectral and frequency-domain analysis. The image sizes of many of the applications have increased over the years reaching 2048 × 2048 in synthetic aperture radar image processing [31], digital holographic imaging [86]. Existing 2D DFT implementations include software solutions, such as FFTW [52], Spiral [117], Intel MKL [3] and IPP [2] which can run on conventional computers, multicore architecture [49], or supercomputers [4]. There are several hardware solutions using the dedicated FFT processor chips [86],[118],[89],[1],[9] and field programmable gate array (FPGA) based implementations [134],[134],[38][83],[99],[132]. These implementations are efficient, however they are incompatible with an asynchronous event based acquisition, as only a fraction of the signal changes over a short time period, therefore applying the FFT on the whole signal would not be efficient and would not make a full use of the advantages of this acquisition process.

State of the art level-crossing sampled signals and more generally stochastic sampled signals (not complying with the Shannon sampling theory) have been studied with the goal of achieving 1D signals accurate reconstructions [133] and frequency analysis (e.g. filtering and Fourier-like transformations) [112, 46, 45, 57]. This work differs from that topic of research because we are dealing for the first time, with the computation of spatial Fourrier transforms of a dynamic scene acquired using an asynchronous event-based image sensor.

2.2 THE ASYNCHRONOUS TIME-BASED IMAGE SENSOR

The ATIS used in this work is a time-domain encoding image sensor with 240 × 304 pixel resolution [115]. The sensor contains an array of fully autonomous pixels that each combines an illuminance change detector circuit and a conditional exposure measurement block. As shown in the functional diagram of an ATIS pixel in Fig. 2, the change detector individually and asynchronously initiates the measurement of an exposure/gray level value only if — and immediately after — a brightness change of a certain magnitude has been detected in the field-of-view of the pixel at time t_0. The ATIS encodes visual information as a stream of events. An event is a set $\{type, \mathbf{p}, t_0, pol\}$: where $type$ is the flag signaling a change, or a gray level event, $\mathbf{p} = (x, y)^T$ the spatial coordinate, pol the polarity, and t_0 the time of occurence. The polarity pol has two meanings according to the event's type. For a change event, the polarity encodes the increase or the decrease of the luminance. For a gray level measurement mechanism it differentiates between the two events encoding the temporal measurment of luminance. Luminance in our case is encoded by a pair of gray level events such that the inverse I of the time difference between the two is proportional to the luminance (as shown in Fig. 2). The linear correspondance between the measured timing and the absolute luminance value is set by design. The first gray level event is triggered right after the change event at $t_1 \sim t_0$ and the second at t_2 such that :

$$I \propto \frac{1}{t_2 - t_0}. \tag{1}$$

Readers are advised to refer to [115] for further details about the ATIS.
Since we are focusing on the luminance information at time t_0, we define a simplified event e_{cam} as:

$$e_{cam}(\mathbf{p}, t_0) = \{I, \mathbf{p}, t_0\} \tag{2}$$

This integration duration only depends on the measured gray level intensity, not on the time t at which it started. Since the ATIS is not clocked like conventional cameras, the timing of events can be conveyed with a temporal resolution of the order of 1 µs. The time-domain encoding of the intensity information automatically optimizes the exposure time separately for each pixel instead of imposing a fixed integration time for the entire array. This results in an exceptionally high dynamic range of 143 dB and an improved signal to noise ratio of 56 dB.

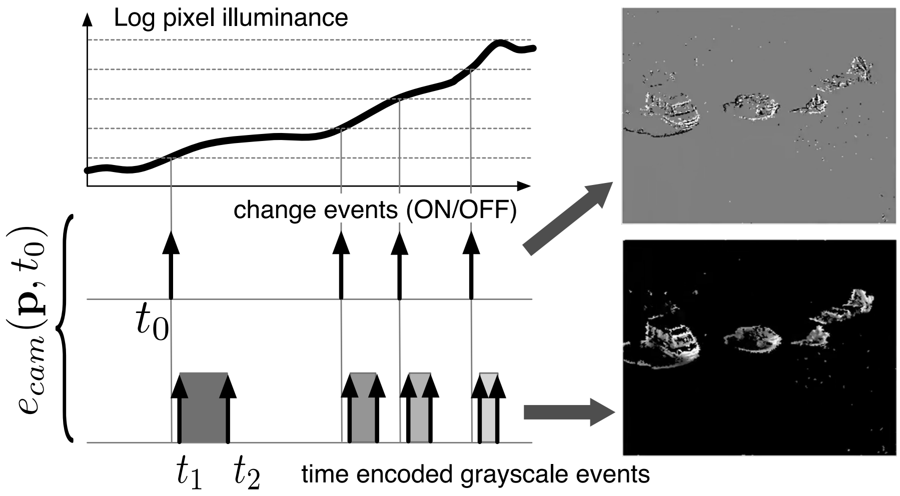

Figure 2: **Functional diagram of an ATIS pixel**. Two types of asynchronous events, encoding change and brightness information, are generated and trans- mitted individually by each pixel in the imaging array.

2.3 DESIGN AND REPRESENTATION OF EVENT-BASED ALGORITHMS

General framework

The event-driven acquisition and the absence of a global sampling frequency radically changes the signal representation and update, compared to the conventional frame-based representation. Only few components of the acquired signal are updated when a change happens. This implies that to benefit from the high temporal accuracy algorithms should be event-driven. This means that processing must only be carried out when an event is acquired in the same spirit to what has been presented in [67]. The whole chain of processing is event driven meaning that an iterative computation must be developed and updated for each incoming event.

The solution proposed in this work is reached in two-stages: a first event-driven naive formulation is built from the standard definition of the DFT. This form, as we will show, is exact but not optimal since a single event updates all Fourier coefficients. A second form is then derived based on the Stockham algorithm [34, 136] using a decomposition into sparse matrices of the DFT operator. This decomposition, combined with the event-based formulation, allows to achieve a lossy DFT that discards non significant events from the DFT computation, hence enabling a strategy based on trade-off between accuracy and computation time.

The naive approach: a simple algorithm to compute the Discrete Fourier Transform

The first step in designing event-based algorithms is to study the impact of sparse event-based updates of an acquired signal and its implications on the computation of the Discrete Fourier Transform. Let's consider a real valued signal $\mathbf{s}(n,t)$ which is a function of space, indexed by n, and of time t. We denote $\mathbf{S}(k,t) = \mathcal{F}[\mathbf{s}(n,t)]$ its spatial DFT computed at t :

$$\begin{array}{rccc} \mathcal{F}: & \mathbb{R}^N & \longrightarrow & \mathbb{C}^N \\ & \mathbf{s}(n,t) & \longmapsto & \mathbf{S}(k,t) = \mathcal{F}[\mathbf{s}(n,t)]. \end{array} \tag{3}$$

$\mathcal{F}$ is the N-point DFT operator with $N \in \mathbb{N}^*$ (set of non null natural numbers) applied to the spatial components of $\mathbf{s}$. To ease understanding, we develop the methodology for spatially unidimensional signals. The same methodology can be extended to multidimensional signal. Experimental results on 2D signals will be presented in section 2.5.

The DFT of $\mathbf{s}$, referred to as the analysis equation, is defined for integer k satisfying $0 \leqslant k \leqslant N-1$ or equivalently $k \in [\![0, N-1]\!]$:

$$\mathbf{S}(k,t) \coloneqq \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} \mathbf{s}(n,t) \exp\left(-2i\pi \frac{nk}{N}\right) \tag{4}$$

and the inverse transform, referred to as the the synthesis equation is:

$$\mathbf{s}(n,t) = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} \mathbf{S}(k,t) \exp\left(+2i\pi \frac{nk}{N}\right) \tag{5}$$

Note that the definition of the DFT is not always the one we present here. We chose this convention for the normalization factor because it enables us to keep the same rules for forward and inverse transforms.

Using this normalization convention, Plancherel's theorem, applied to the spatial component of $\mathbf{s}$ at a given t and to its spatial Fourier transform, is:

$$\forall t, \|\mathbf{s}(n,t)\|_2 = \|\mathbf{S}(k,t)\|_2, \tag{6}$$

where $\|.\|_2$ is the Euclidean norm in each respective space. Because of that property, the same threshold, referred to as the "significance threshold T", can be used both in the focal plane and in the frequency space. This thresholding mechanism is the basis of the idea developed in section 2.4.

When an event $e_{cam}(q, t)$ is detected and acquired it implies that a single component of the input signal s at location q changes significantly. Here, we are considering a 1D case to ease the notation, therefore q is a scalar representing the index corresponding to the location where the change occurs. We denote $\mathbf{s}(n, t)$ the acquired signal, at time t. In order to disambiguate the value of the acquired signal at the event times, we use the following notations :

$$\forall p \in [\![0, N-1]\!], \begin{cases} s(p, t^-) & := & \lim\limits_{u \to t^-} s(p, u) \\ s(p, t^+) & := & \lim\limits_{u \to t^+} s(p, u) \end{cases}$$

The acquired input signal is then updated at its component q, so that $e_{cam}(q, t) = \{\mathbf{s}(q, t^+), q, t\}$. The values of the acquired signal right before and right after the event $e_{cam}(q, t)$ are related by :

$$\forall p \in [\![0, N-1]\!], \mathbf{s}(p, t^+) = \mathbf{s}(p, t^-) + \alpha\delta_{p,q} \tag{7}$$

where δ is the Kronecker delta and α is the difference between the old value and the new value of the signal.
We can establish the relation between $\mathbf{S}(k, t^+) := \mathcal{F}[\mathbf{s}(n, t^+)]$ and $\mathbf{S}(k, t^-) := \mathcal{F}[\mathbf{s}(n, t^-)]$:
$\forall k \in [\![0, N-1]\!]$,

$$\begin{aligned} \mathbf{S}(k, t^+) &= \tfrac{1}{\sqrt{N}} \sum_{p=0}^{N-1} \mathbf{s}(p, t^+) \exp\left(-2i\pi \tfrac{pk}{N}\right) \\ &= \tfrac{1}{\sqrt{N}} \sum_{p=0}^{N-1} [\mathbf{s}(p, t^-) + \alpha\delta_{p,q}] \exp\left(-2i\pi \tfrac{pk}{N}\right) \\ &= \tfrac{1}{\sqrt{N}} \sum_{p=0}^{N-1} \mathbf{s}(p, t^-) \exp\left(-2i\pi \tfrac{pk}{N}\right) \\ &\qquad\qquad + \tfrac{\alpha}{\sqrt{N}} \exp\left(-2i\pi \tfrac{qk}{N}\right) \\ &= \mathbf{S}(k, t^-) + \tfrac{\alpha}{\sqrt{N}} \exp\left(-2i\pi \tfrac{qk}{N}\right), \end{aligned} \tag{8}$$

This is showing that every term of the DFT has to be updated with an increment with the same module $|\alpha| / \sqrt{N}$, and a phase which depends on the indices of both the pixel component and the Fourier component. Consequently, updating the exact Fourier spectrum after one event on the camera requires a number of operations linear with the number N of samples. It is straightforward to conclude that no exact iterative method can be implemented in less than $\mathcal{O}(N)$ operations per event.

Finally, it suggests that no approximation can be made *a priori*. Since all components should be updated by an increment of the same amplitude, we can not know which ones to favor and which ones to exclude (except possibly for a specific application for which a fraction of components are more important than others, which is not the case

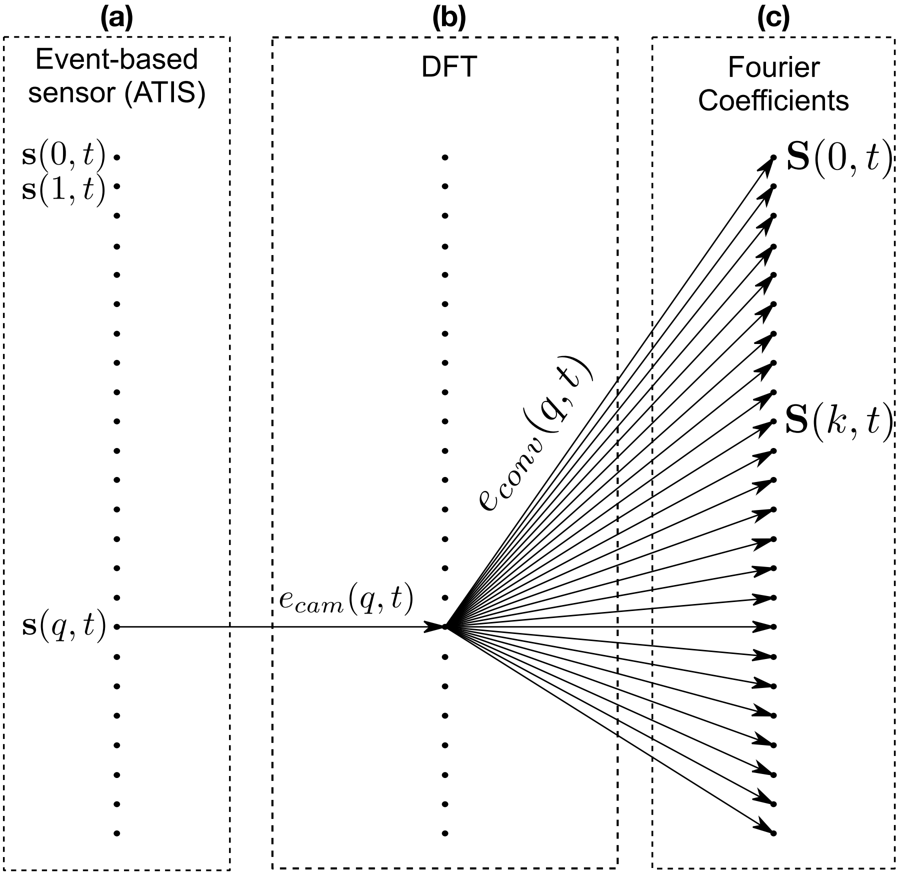

Figure 3: **Graphical representation of the naive event-based algorithm** (unidimensional case, $N = 24$). (a) the gray level events are generated from the ATIS. They are connected in a one-to-one manner to the converter (b) which computes the difference between two successive events. The nodes of the converter implement the computation of variable I of (7). The Converter layer is connected in an all-to-all fashion to the Fourier layer (c). The Fourier nodes implement the computation introduced in (8).

here as we intend to provide a general method).

The algorithm consists in three successive blocks as represented in Fig. 3.

The sensor is delivering events $e_{cam}(q,t)$ to each node of the DCT block. Each node then uses its index $k \in [\![0, N-1]\!]$ to compute the value of the update, $\frac{\alpha}{\sqrt{N}}\exp\left(-2i\pi\frac{qk}{N}\right)$ that is is added to the previous value of the k^{th} Fourier component $\mathbf{S}(k,t^-)$ to compute the new value of the component $\mathbf{S}(k,t^+)$. This value replaces the previous one in memory and an is output by that third layer. It is important to notice that the rate of events (i.e. the number of events triggered/processed per unit of time by each layer) is multiplied by N at this stage, where N is the number of spatial samples of the input signal.

Remarks :
The algorithm has two main limitations. The first one is that the number of operations carried out per event is linear with the number of pixel of the sensing device. As an example, the ATIS camera has $240 \times 304 = 72960$ pixels. Performing that much operations per event is too resource demanding and does not make full use of the event-driven properties of the sensor.

The second limitation is that it also increases the rate of generated events by a factor 72960, because for each event output by the ATIS, each node representing a Fourier component triggers an event containing its new value.

It is important to emphasize that the terminology of "event" applies for all the entire processing chain, there are two types of events, those provided by the ATIS are actual measurements of the absolute luminance, the remaining ones are purely computational output by each layer.

2.4 EVENT-BASED DISCRETE FOURIER TRANSFORM

A better methodology is to use a decomposition of the matrix representing the DFT operator into a product of sparse matrices. We can apply Stockham's algorithm [34, 136] that computes the DFT in $\mathcal{O}(N \times \ln(N))$ operations, and has both the input and output signals sorted in the natural order.

The computation of the DFT using this technique is equivalent to building a network consisting of several layers, each containing the same number of nodes which is also the number of samples of the input signal and storing the result of intermediate calculations. If we denote M the matrix representing the DFT operator (i.e. $\mathcal{F}[\mathbf{s}(.,t)] = M\mathbf{s}(.,t)$), FFT algorithms provide us with decompositions of the matrix M as a product of L sparse matrices, where L is of the order of $\ln(N)$:

$$M = M^L \dots M^2 M^1 \tag{9}$$

Given an input signal which is stored in the nodes of the first layer, the values of the nodes of the following layers are equal to a weighted sum of the nodes in the previous layer which connect to it. A weight is associated with each edge. Computation of the connections and associated weights within the network can be performed based on the knowledge of the sequence of factors chosen to build the network. The signal contained in each layer is obtained as a linear combination of the signal contained in the previous layer. A network is equivalent to a sequence of matrices, where each matrix is associated with the transition from one layer to the next. We can compute the number of operations — defined as a multiplication of a complex by a complex exponential followed by a complex addition — such propagation requires. Each edge accounts for one such operation. Fig. 4 shows such a network for a unidimensional input signal, with $N = 24$. It is built using the decomposition $24 = 2\times2\times2\times3$, an update would require $2+4+8+24 = 38$ operations. This is an expected result as updating exactly the Fourier representation requires at least N operations. Unsurprinsingly, there is a path from each node of the input layer to all nodes of the output layer.

To make full use of the event-driven acquisition we ensure that an incoming event introduces a significant change to the signal. Instead of sending an event each time a value is updated, events are only propagated if the update they communicate is significant. As shown in Fig. 4 an incoming event from the event-based camera $e_{cam}(20, t^+)$ from node 20 is sent to the layer 1 of the DFT block, because the amount of change it provides is larger than a percentage of the previously received value of the $\mathbf{s}(20, t^-)$. Otherwise this new information is stored locally by updating the value of the node until the amount of the change is significant enough with respect to a fixed threshold T.

A example of successful propagations of the signal changes within the DFT block can be depicted as follows:

- The ATIS outputs an event $e_{cam}(20, t^+)$ at its node 20. This event is transferred to nodes 8 and 20 of the first layer of the DFT block for its first layer processing.

- Out of the new values at nodes 8 and 20 in the 1st layer, only the one for node 20 is supposed to be significant enough. This triggers the event $e_{1,2}(20, t^+)$ that is sent to nodes 8 and 20 of the layer 2 for the 2nd layer processing.

- Now only change at node 8 of layer 2 is supposed to be significant. Due to the same mechanism, event $e_{2,3}(8, t^+)$ is generated at node 8 and transmitted to nodes 5 and 17 of the last layer for processing.

- Finally, the new values at nodes 5 and 17 of the last layer are significantly larger that the previous stored values, this triggers events $e_{3,S}(5, t^+)$ and $e_{3,S}(17, t^+)$ that are updating some Fourier Coefficients $\mathbf{S}(k, t^+)$.

All connections corresponding to a full computation are shown in dashed lines (for the last block providing Fourier coefficients, these connections are not shown to emphasize the sparsity of the event-based DFT algorithm and to preserve readability). The event-based thresholding optimization pathway of information triggered by a single incoming event is displayed as plain arrows.

This process introduces an approximation in the computation of the DFT, since the output of the Fourier layer is not the exact DFT of the acquired signal but a signal which is considered to be so close that the difference between the exact and the approximate signals is not worth communicating. Fig. 5 shows an experimental distribution of the number of Fourier components for an outdoors urban complex scene which are updated in a significant manner each time the input signal is updated significantly for a local change of 1%. The number of operations saved using this heuristic algorithm will be measured practically in the experiments section, as expected there will be a tradeoff of quality versus number of operations imposed by the chosen threshold.

2.5 EXPERIMENTS

Methods

Implementation

Recorded sequences of events from the event-based cameras are used. We used the maximum number of layers by decomposing height and the width of the frame into prime numbers :

$$\begin{cases} 240 &= 2^4 \times 3 \times 5 \\ 304 &= 19 \times 2^4 \times 1 \end{cases}$$

Image comparison

The quality of transforms is assessed using the Mean Structural SIMilarity index (MSSIM) introduced in [139] computed every 10 ms. The Structural SIMilarity (SSIM) index is a full-reference measure of similarity between two thumbnail images of 11×11 pixels. Larger images are compared by building 11×11 neighborhoods centered in each pixel. The Mean of the SSIM value is computed over all these possible neighborhoods in the image. This method is suited to measure image distortions, hence it is used to assess to which extent our method is able to maintain the structure of the input signal with respect to the approximations made.

The index is a combination of three measurements : (i) a luminance similarity index, (ii) a contrast similarity index and (iii) a structure similarity index based on the normalized correlation between the two thumbnails. The values for the different parameters were all set using the values recommended in [139]. The index in this DFT context is:

$$\text{MSSIM}(t) = \text{MSSIM}\big(\mathbf{s}(n,t), \tilde{\mathfrak{F}}^{-1} \circ \tilde{\mathfrak{F}}[\mathbf{s}(n,t)]\big), \tag{10}$$

where $\tilde{\mathfrak{F}}$ is the approximation of the DFT operator resulting from our algorithm.

Computation time

Conventional frame-based FFT algorithms take advantage of the structure of the Fourier basis such as the symmetries and periodicities of the trigonometric functions while event-based acquisition implies that the number of operations depends on the the signal. Consequently, our assessment of the results will be experimental based on the statistics of recorded scenes. We ensured a wide variety of indoor and outdoor scenes. We recorded the number of operations required to process each acquired sequence for a range of threshold over the relative change of the signal values : 0%, 1%, 2%, 5%, 10% and 20%.

As shown in section 2.4 (compare Fig. 3 and Fig. 4) the approximation algorithm introduces additional operations when computing the exact Fourier transform. Consequently, our goal is to be more efficient than the naive event-based and the conventional frame-based FFT (when dealing with high frame rates) when using the heuristic approach in terms of computational time.

The number of operation for the frame-based DFT is proportional to the frame rate while the event-based DFT is proportional to the number of measured events. There is no direct and obvious relation between the two numbers. Our only way to determine which of the DFT techniques is requiring the less number of operations for a given quality of reconstruction is to test for different values of T, the DFT algorithms applied to one of our sequences. The comparison is scene dependent as it is shown in the *Results* section, with the moving vehicle sequence.

Results

We present the results obtained for indoor and outdoor scenes shown in Fig. 6. We first consider the most complex recorded scene corresponding to a dynamic urban scene where the camera in mounted inside a moving vehicle (the first row of Fig. 6).

Results of recomposing the output after the heuristic approach for different significance threshold values are shown in Fig. 7. Up to a threshold of 5% of the dynamic range, the structure of the image is very well preserved as well as the details of the image. For a threshold of 10%, the structure is still preserved, but most of the details are lost. Finally, for thresholds higher than 20%, the structure of the resulting image is distorted. Only large objects are recognizable, but their details are lost and their shapes are also distorted. In particular, for a threshold above 10%, it appears that the spatial position of the van is slightly delayed in the image, there is a latency in the update of the spatial position due to the large threshold value. The higher the threshold, the less new incoming events will update the FFT, hence the delay can be large when we are comparing the output with frames generated at the same time. However the rate of events is also largely scene dependent (multi targets entering the sensor field of view, change of the relative speeds,...) and it can impact significantly the delay. At that stage, there is no straightforward way to keep tract of the delay w.r.t. the scene, hence we are evaluating the performance of the approach in the least favorable case where we are not taking the latency into account. Fig. 8 provides the evolution the similarity index for all thresholds for the sequence. It can be noticed that the similarity index value degrades to 0 as the threshold increases to 20%.

Results plot in Fig. 9 for all acquired sequences show that the quality index, MSSIM, and the average number of operation per event A, applied to the data are independant from the scenes' content. The threshold T is the parameter that sets the compromise between the signal transform quality and the computation used for the transform. Significant computational gains are obtained for threshold values in a range between 2.5% and 5%. As expected, the MSSIM and A functions have the same behaviour with respect to the threshold: they are decreasing functions of T. The proposed algorithm takes advantage of the fact that the MSSIM decreases at a much slower rate than A when T increases. Fig. 9) that the the tangent to the MSSIM at zero is almost flat, while the slope of the tangent to A at zero is steep. This behavior allows us to find a threshold such that computations are significantly reduced with a low loss in signal quality. The combination of both functions allows to find a trade-off between the computation time and the quality of the transformed signal, as shown in Fig. 10.

The PSNR (in dB) and the MSE (in gray level amplitude squared) of the reconstructed signal are also plotted in Fig. 11 to show the impact of the threshold T. The gray levels measured by the ATIS are normalized by the highest value and rescaled so the gray levels values are between 0 and 255. These curves are substantiating the conclusion drawn from the MSSIM: the reconstructed signal quality is degraded quickly as T increases. For $T = 5\%$, the threshold value for which the tradeoff between MSSIM measured quality and computation time is becoming less interesting, the PNSR and the MSE are respectively 22.7dB and 346.

To compare the frame-based Fourier transform with the algorithm we introduced, we are estimating the number of operations per unit of time. This quantity is obtained by multiplying the average number of operations performed per event $A(T)$ by the number of events per unit of time. This value is estimated experimentally for the moving vehicle sequence and is equal to 10^6 events per second. Fig. 12 shows the rate of operations w.r.t. T as a decreasing function when the sensor is static (blue dots) and when it is put in a car (red dots). Three dashed lines are plotted to show the number of operations used by the frame-based algorithm at three different video sampling frequencies ($100\,\text{Hz}$, $1\,\text{kHz}$ and $1\,\text{MHz}$). At $1\,\text{kHz}$, for $T = 6\%$, both techniques require the same amount of operations. This amount decreases even more for the event-based algorithm if T increases.

As shown in Fig. 12 and Fig. 7, threshold values from 6 to 8% of the dynamic range are a compromise for computing with reasonable resource, the Fourier transform with respect to the loss of quality.

2.6 DISCUSSION

As shown in the experimental section, the update process in a series of steps leads to significant gains in computational time. These gains conventionally imply a loss in accuracy because each intermediate layer filters out events that lead to small increments in all following steps. Interesting gains occur at low threshold values (less than 1% of the dynamic range) where significant low computational time are achieved while virtually no change can be detected in the signal. This shows that the heuristic algorithm determines which components should be updated. The algorithm does not require *a priori* choices regarding the components which should be updated, but rather bases the decision on the incoming signal. As the algorithm is intended to be applied to a wide range of areas which make use of level-crossing sampling, or more generally of asynchronous sampling, it was our choice not to make use of any prior knowledge in the design of the architecture of the system.

However, the algorithm could be adapted to benefit from prior knowledge. Two promising leads could be (i) adapting the thresholds so that more computational resources are allocated to components of higher interest, and (ii) re-arranging the connections in the network in order to filter out events in the earliest possible layers. Regarding the second point, the decomposition we used here is based on Fast Fourier Transforms decompositions of the Fourier operator. These decompositions result in minimal numbers of connections for a given number of samples and a given number of intermediate layers, which is a sensible approach to start with. However, neighboring pixels, which are the most likely to be related in a visual signal, do not connect to neighboring nodes in the first intermediate layer (see Fig. 4).

We provided a trade-off (see Fig. 10) as a characterization of our algorithm, and did not look for any optimal value for the threshold. Such value for the threshold heavily depends on the application for which it is used. A specific application would provide an objective function for both the computation time and the accuracy of expected results. The trade-off curve, or a similar one which would be produced for the given application, would then allow to turn the objective function of accuracy and computational time into a function of the threshold. Maximizing the resulting function with respect to the threshold would then provide the optimal threshold.

Considering the heuristic method, the perfect algorithm would provide a network and a behaviour for the nodes such that (i) the number of layers scales with $\log(N)$ and (ii) the rate of events through each layer is kept constant on average. In such scenario, the number of operations carried out per event would scale with $\log(N)$, and the corresponding algorithm would lead to the same improvement as the one provided by FFT algorithms in the frame-based setting. The fact that we filter events out suggests that the approximation method scales between $\mathcal{O}(\log(N))$ and $\mathcal{O}(N)$.

Finally, increasing the threshold does not only allow computational gains within the Fourier filter, but it also decreases the rate of events output by the filter. This in turn reduces the computational burden of further processing steps.

2.7 CONCLUSION

In this paper, we provided an algorithm to implement event-based Fourier transform algorithms. As the demand for higher temporal resolution increases, in particular for artificial visual tasks, the need for update methods which are able to operate on sparse data representation will be increasingly high. We showed that a promising lead is to develop heuristics which are able to regroup incoming information in order to detect as soon as possible, i.e. in the earliest possible layer, which part of the information is unnecessary to propagate to the following steps. As in the different frame-based FFT approaches, important work can be carried out by comparing decompositions of the Fourier operator matrix. However, the event-based framework introduces a major shift in the objective of the decomposition. The goal is not to find a decomposition which provides minimal number of edges in the associated network, but rather its ability to filter out useless information which in turn depends heavily on the statistical structure of the input signal, and of its dynamic. Our work introduces a general framework showing that adding intermediate computation steps can help reducing the computational burden with minimal degradation of the underlying signal.

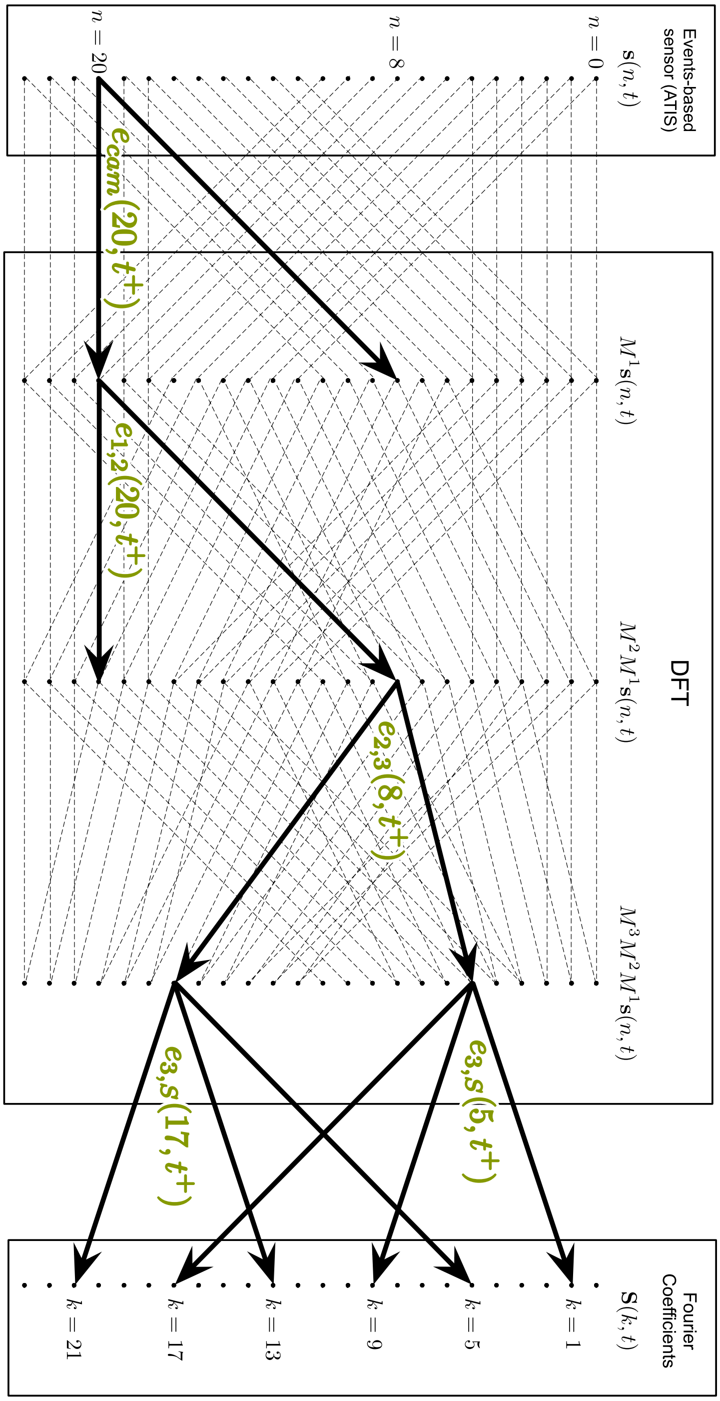

Figure 4: **A graph used for the computation of a unidimensional dynamic signal with 24 samples**. The values of the input signal are stored in the leftmost nodes. Edges show the relationship between the values contained in the nodes of the network. The value contained in a node is a linear combination of the values contained in the nodes of the previous layer (i.e. to the left) connected to the node through an edge. The weights in the linear combination depends on the edge, but are omitted in the figure. The solid edges show the path of an update from input node number 20 to all output nodes.

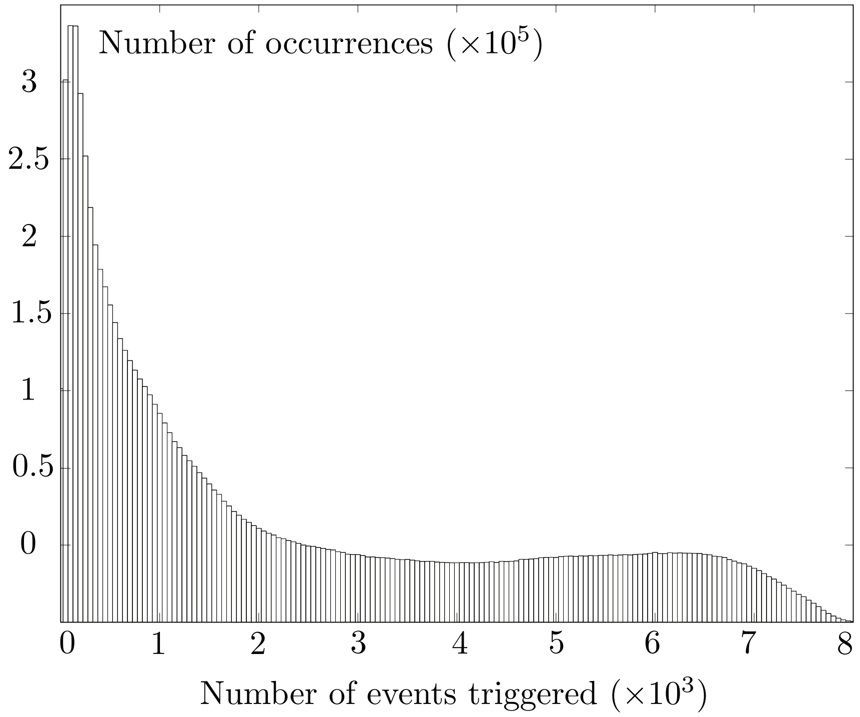

Figure 5: **Histogram of updated Fourier coefficients from the naive DFT**. We performed the event-based Fourier algorithm on a set of natural scenes videos recorded with the ATIS. For each input event representing an update of more than 1% of the dynamic range, we counted the number of Fourier components which were updated by an increment superior to the same threshold. The worse condition happens when all the Fourier coefficients are updated despite of the significance thresholding, in that case, the number of updated coefficient is equal to the number of pixels in the sensor i.e. 72960.

Figure 6: Snapshots from the 8 event-based sequences used for the experiments.

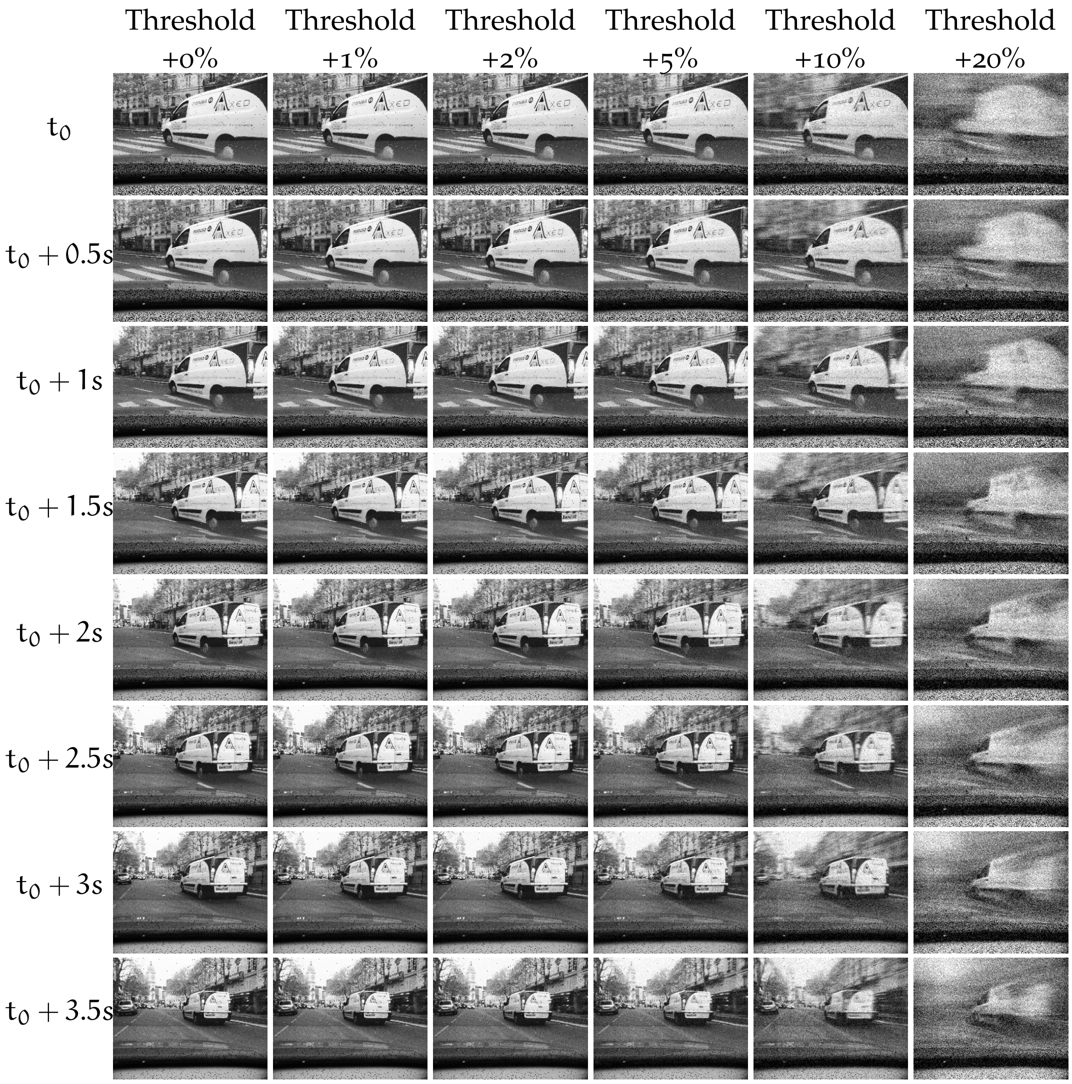

Figure 7: ***Moving vehicle* sequence**: signals obtained after computing the Discrete Fourier Transform and its inverse using the algorithm presented in 2.4. The same sequence (leftmost column : Threshold 0%) is processed using different values for the significance threshold T ranging from 1% to 20% of the dynamic range. Images are reconstructed using the flow of gray level events every half second.

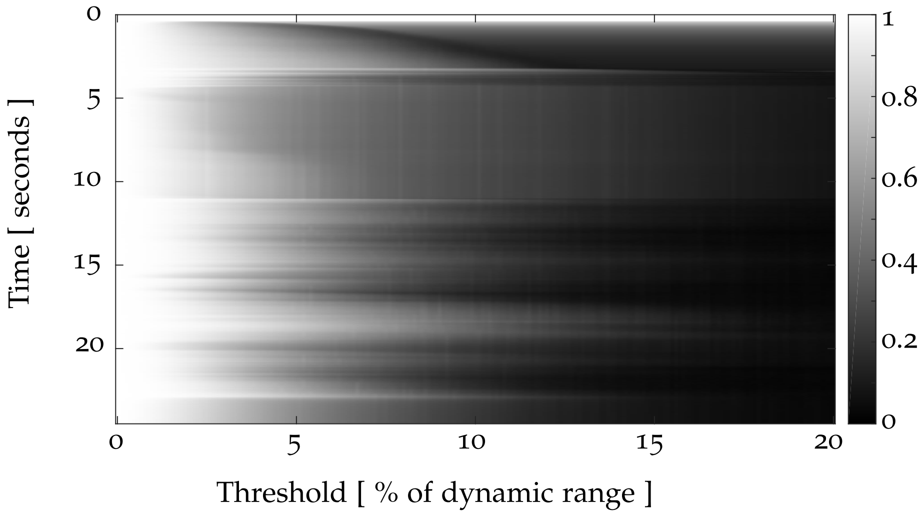

Figure 8: Variations of the Mean SSIM index for the dynamic urban sequence, w.r.t. the threshold (in % of the dynamic range) and time.

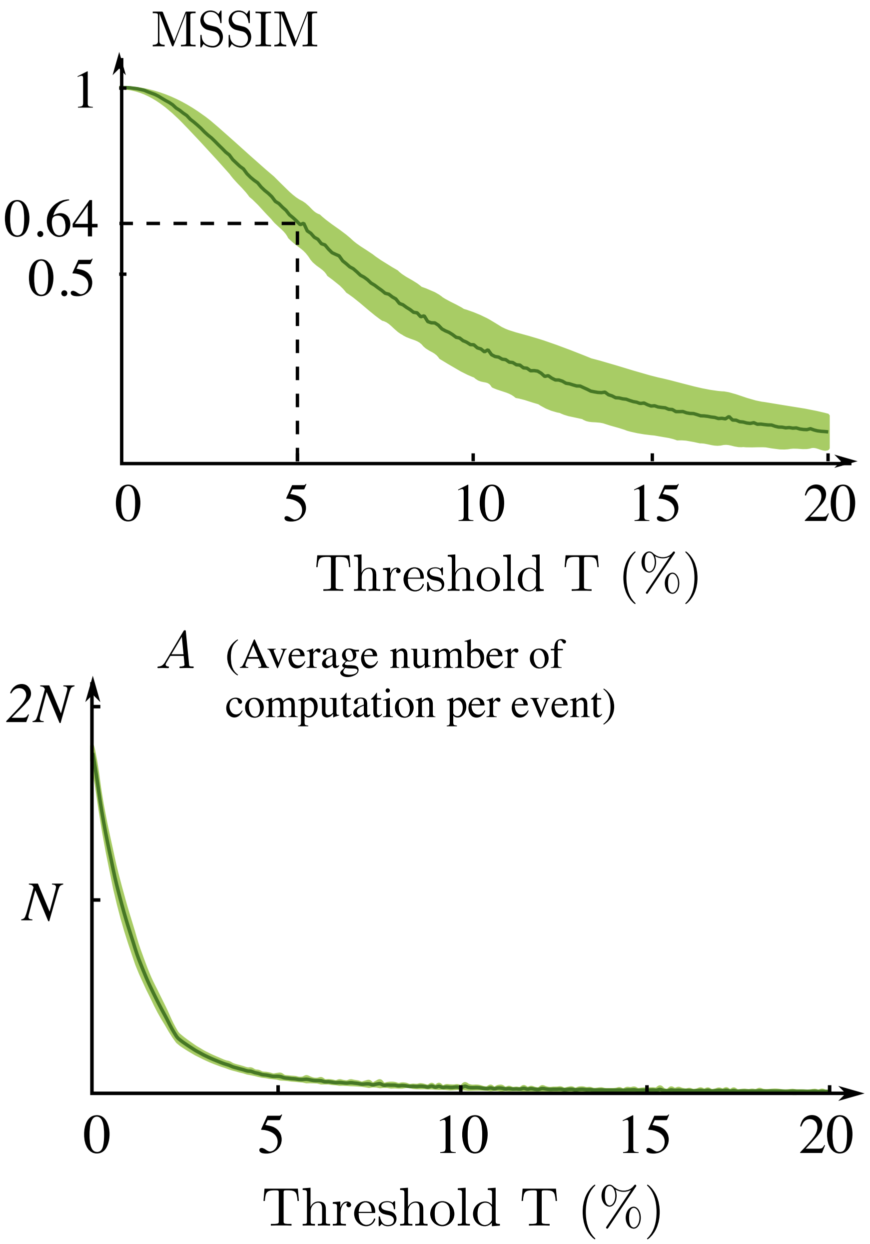

Figure 9: MSSIM and average number of operations per events are decreasing functions of the threshold. One can however see the MSSIM is decreasing in a much slower rate.

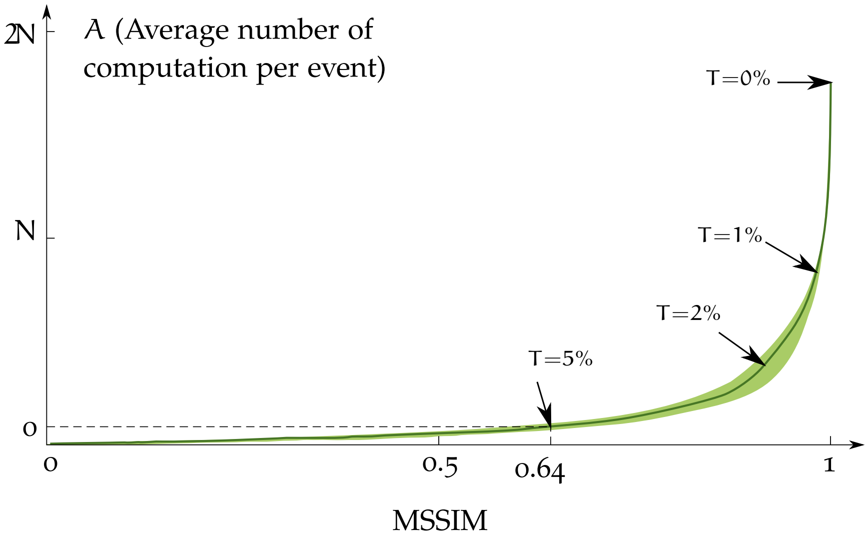

Figure 10: Average number of operation per events w.r.t. the quality index. This graph emphasizes the slow increase rate of A when the MSSIM increases.

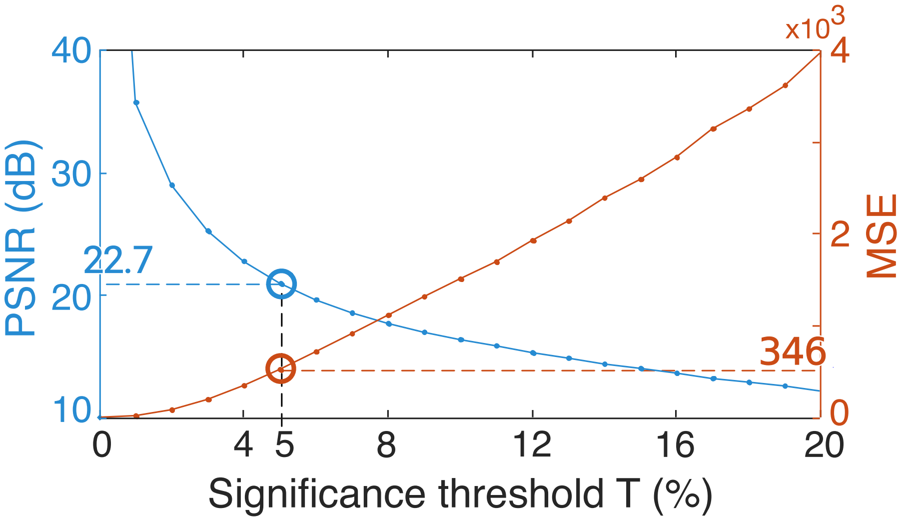

Figure 11: PSNR and MSE as function of T. The PSNR decreases in a similar fashion as the MSSIM which is reflected by the increasing behavior of the MSE.

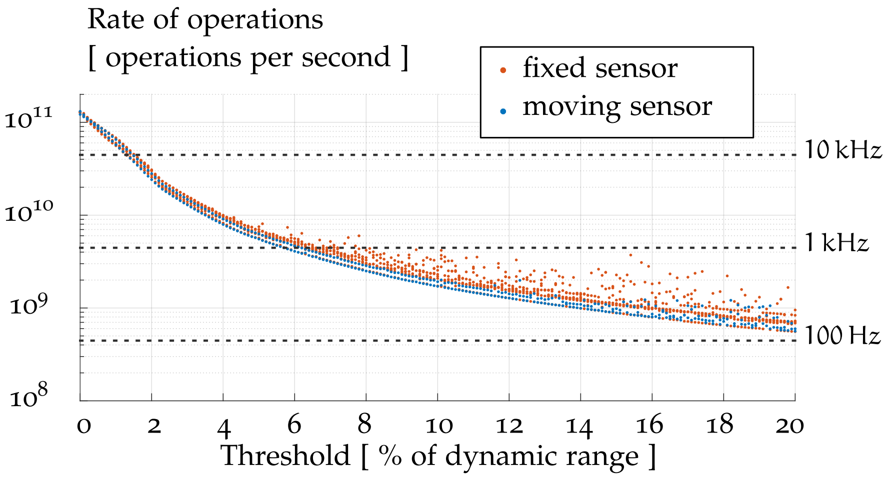

Figure 12: **Number of operations per unit of time carried out during our algorithm** assuming a mean event rate of 10^6 events per second. Horizontal dashed lines show the rates of operation required by the frame-based FFT algorithm at different sampling frequencies : 100 Hz, 1 kHz and 10 kHz. The rate for the event-based algorithm is larger than the frame-based algorithm for low sampling frequencies up to 100 Hz. When sampling frequencies go up to the kHz, which corresponds to the typical µs precision of the asynchronous sensor, the event-based algorithm is more computation efficient for threshold values around 6-8% of the dynamic range.

Part III

MODELING THE RESPONSE OF THE TRANSFECTED CELL

3

MODELING THE BEHAVIOR OF CHRIMSONR

Channelrhodopsins are light-gated cation channels which are naturally expressed in microalgae. This guides the *photoaxis*, the ability for the microalga to move towards optimal light conditions[62]. They belong to the larger family of type I rhodopsins, along with halorhodopsins and bacteriorhodopsins, and mediate vision in microorganisms.

In 2002, Nagel and colleagues reported the DNA sequence encoding a microbial opsin Channelrhodopsin-1 (ChR-1) expressed in the green alga *Chlamydomonas reinhardtii*[101]. They managed to express the protein in another organism, the *Xenopus laevis* oocyte, and observed its electrochemical behavior under light stimulation. One year later, they reported the discovery of another channelrhodopsin from the same green alga, termed Channelrhodospin-2 (ChR-2); identified as a cation-selective light-gated channel[102].

A major step in the development of optogenetics was the expression of ChR-2 into mammalian neurons and its use in controlling the spiking behavior of transfected neurons[24]. Contrary to other type I rhodopsins, channelrhodopsins depolarize neurons in which they are expressed and induce action potentials when activated. Additionally, it showed that the retinal chromophore was covalently bound to the apoprotein, and thus that no supplementation of retinal to the culture medium was needed. This started the field of optogenetics. Since then, it has been used to control neuronal activity of specific cell populations (e.g. in [5, 7]).

In parallel, an important effort has been devoted to understand the mechanism underlying the behavior of channelrhodopsins. Studies include (i) spectroscopy to determine the relevant conformations of the protein during light excitation, (ii) crystallography which reveals the precise structure of the protein and location of the amino acid residues and retinal chromophore[100, 74], (iii) electrophysiology to study the dynamics of the conductance induced by light stimulation[65, 80, 127, 94], (iv) numerical simulations[63, 105, 47] and (v) site-directed mutagenesis[88, 15, 81, 59, 141] targeting putative key residues to test their influence on experimental observations and infer their role in the photocycle.

It led to establishing the outline of the photocycle[14, 119]. Starting from the dark-adapted conformation termed D470, and upon photo-induced isomerization of the retinal chromophore, the protein takes a number of intermediate states P500, P390 leading to the conductive state P520 (see Fig. 13). From this conductive state, the protein either relaxes to the initial dark state or to another closed state P480, which is thought to be photoactivable, and which relaxes slowly (approx. 20 seconds)[93]. In the same time, site-directed mutagenesis allowed to identify the residues most likely involved in the function of the light-gated channel, for instance as proton donor or acceptor allowing the Retinal Schiff Base deprotonation and reprotonation[93, 81].

However, voltage-clamp experiments reporting the conductance response of the channels to light stimulation showed that the observations could not be explained by a single-loop photocycle, contradicting the preliminary conclusions from spectroscopy studies. In particular, when light is switched off, the conductance returns to the baseline according to a linear combination of two exponential terms, with two distinct time constants[63, 105], phenomenon which requires at least two conductive states. Additionally, voltage-clamp experiments at different holding potentials showed that the cation-selectivity of the two open states were different, as shown by the distortion of the conductance responses for different holding potentials[19, 127].

It is assumed that the two photocycles have the same shape, each one consisting of twin states which are not distinguishable from a spectroscopic point of view[130, 15, 121, 80]. However, it remains unclear where the transitions between the two cycles occur, and whether they are thermal or light-induced. It was first proposed that the transitions occur between the inactivated states P480 and P480'[121]. Recently, it was shown that a light-induced transition existed between the two dark adpated states D470 and D480[26].

Additionally, the possible transitions within a single cycle are not fully elucidated. In fact, since the transition P480 → D470 appears to be very slow (e.g. 20 seconds in [94]) another transition has to be added in the cycle — either a thermal transition from P520 to D470 as proposed in [94] or a light-induced one from P480 to P500 as in [18] — otherwise, the electrophysiology studies would not show a sustained response to continuous light stimulation.

Finally, a side reaction was revealed in [130] on the photocycle of the ChR-2 mutant C128T. This reaction occurs between states P480 and D470 and involves the inactivated states P380 and P353 that convert very slowly back to D470. This side reaction was placed in the two-photocycle model[121].

Note that electrophysiological experiments have been conducted on different cell types including *Xenopus* oocytes, human embryonic kidney (HEK) cells, baby hamster kidney (BHK) cells, Henrietta Lacks (HeLa) cells, cultured neurons, etc... and their properties have been shown to be mostly insensitive to the host system used[94, 105].

This paper focuses on ChrimsonR, another channelrhodopsin which was discovered upon a vast scanning of the transcriptomes of more than 100 species of alga[78]. ChrimsonR is 45 nm more red-shifted than all previously known channelrhodopsins. Its peak sensitivity is at 590 nm. Sensitivity to this wavelength is valuable for three reasons it is more suitable for medical applications; it allows efficient stimulation in deeper tissues; it can be used for independent stimulation of two neuronal populations when used jointly with another channelrhodopsin with a non-overlapping activation spectrum, like Chronos[78].

Based on voltage-clamp experiments on ChrimsonR-expressing HEK293 cells, we report new observations on the response behavior of channelrhodopsins . Additionally, we propose a five-state Markov kinetic model that captures most of the features that we describe : the on and off dynamics, the activation curve (steady-state amplitude of the response as a function of light intensity). In the end, our model is able to predict accurately the response of the protein on long time scales (several minutes) for temporal frequencies lower than approx. 100 Hz and over the whole range of light intensities relevant to the activation of the protein (between no response at 10^{15} ph s^{-1} cm^{-2} and saturation around 10^{19} ph s^{-1} cm^{-2}).

The scope of this work is to describe relevant kinetics and expected behavior of the channel. This will help design appropriate light-stimulation protocols for optogenetics-based vision restoration.

3.1 METHODS

HEK 293T cell culture, transfection

HEK 293T cells were maintained between 10% and 70% confluence in DMEM medium (Invitrogen, Waltham, USA) supplemented with 10% FBS (Invitrogen), 1% penicillin/streptomycin (Invitrogen). For recording, cells were platted at 50,000 cells per well in 24-well plates that contained round glass coverslips (12 mm) coated with polylysine (2 μg cm^{-2}, Sigma Aldrich) and laminin (1 μg cm^{-2}, Sigma Aldrich). Adherent cells were transfected approximately 24h post-plating with JetPrime (Polyplus Transfection) and recorded via whole-cell patch clamp between 24 and 72h post transfection. 0.5 μg of DNA was delivered per well. In addition, all-trans retinal (ATR, 10 μM) was supplemented to the culture medium for 1h before patch-clamp experiments.

The opsin ChrimsonR was expressed into the HEK cells associated to a fluorescent protein tdTomato fused to its C-terminal end. Transfection was performed through the mutant capsid AV2-7m8[39]. This construct was chosen for its efficiency in transfecting retinal ganglion cells through intravitreal injection.

Electrophysiolgy

Whole-cell patch-clamp recordings were performed in isolated HEK 293T cells to avoid space clamp issues. All recordings were performed using an Axopatch 200B amplifier and Digidata 1440 digitizer (Molecular Devices) at room temperature. Steady access resistance were inferior to 40 $\mathrm{M\Omega}$. Typical membrane resistance was between 200 $\mathrm{M\Omega}$ and 1 $\mathrm{G\Omega}$, and pipette resistance was between 5 and 8 $\mathrm{M\Omega}$. Cells were perfused with Ames medium (Sigma-Aldrich, St Louis, MO; A1420) bubbled with 95% O_2 and 5% CO_2 at 37 °C at a rate of 1-2 $\mathrm{ml\,min^{-1}}$ during experiments. Intracellular solution was composed of (in mM) 115 K-gluconate, 10 KCl, 0.5 $CaCl_2$, 1 $MgCl_2$, 1.5 EGTA, 10 HEPES, 4 ATP-Na_2, pH 7.3 (KOH adjusted).

Illumination

Photostimulation of patch-clamped cells was conducted with a 595 nm LED (M595L3, Thorlabs, halfbandwidth of 75 nm). Irradiance was measured, for different LED voltages, at the level of the coverslip sample using a power meter composed of a digital optical power and energy meter (Thorlabs, PM100D) and a photodiode power sensor (Thorlabs, S120C). Resulting power measurements were converted in $\mathrm{ph\,s^{-1}\,cm^{-2}}$ using known illuminated surface (1.21 $\mathrm{mm^2}$) and peak wavelength for the LED (595 nm). These measurements were confirmed using a calibrated spectrophotometer (USB2000+, Ocean Optics, in-house calibration).

Illumination protocols

The building block of our stimulation protocol is a series of ten 200 ms square pulses repeated at a frequency of 0.5 Hz. These building blocks were then assembled in several ways with different light levels. The duration between two subsequent blocks was not set *a priori* (as it required replacing manually an optical filter between the light source and the recorded cell) but the actual value was recorded and usually lasted about ten seconds.

Light intensity of stimulation blocks covered the whole range of responses from the protein, from virtually no response at 3×10^{15} $\mathrm{ph\,s^{-1}\,cm^{-2}}$ to saturation in amplitude at 10^{19} $\mathrm{ph\,s^{-1}\,cm^{-2}}$.

Data analysis and simulation

Markov kinetic models The behavior of the protein ChrimsonR is modeled using a Markov kinetic model. In this model, a number of *states* represent the different conformations that the protein can take. For each pair of states, there can be a transition between these two states if the protein can switch from the first state to the other without going through any other stable conformation. It is worth noting that transitions are not necessarily reversible, i.e. if the transition from state i to state j exists, the transition from j to i does not always exist. Additionally, a time constant is associated with each transition. For a pair of states (i, j), the rate $\lambda_{i,j}$ — the inverse of the time constant $T_{i,j}$ — quantifies the probability that the protein jumps from state i to state j per unit of time. If there is no transition from state i to state j, this transition has a zero transition rate, or equivalently an infinite time constant.

For a given model with N states, and a transition matrix $\Lambda = (\lambda_{i,j})_{1 \leqslant i,j \leqslant N} \in \mathcal{M}_{N,N}(\mathbb{R}^+)$, and denoting $p_i(t)$ the probability that the protein is in state i at time t, the evolution of the system is given by $P(t) = (p_i(t))_{1 \leqslant i \leqslant N}$ and follows the system :
$\forall i \in [\![1, N]\!]$

$$\frac{dp_i}{dt}(t) = -\left(\sum_{j \neq i} \lambda_{i,j}\right) p_i(t) + \left(\sum_{j \neq i} \lambda_{j,i} p_j(t)\right) \tag{11}$$

In the sequel of the paper, we will consider two types of transitions, thermal and photo-induced. Thermal reactions involve absorption or evolution of heat. They can take place even in absence of light. Temperature has a significant effect on the rate of a thermochemical reaction. On the other hand, photochemical reactions involve absorption of light. The presence of light is the primary requisite for the reaction to take place. Temperature has very little effect on the rate of a photochemical reaction. Instead, the intensity of light has a marked effect on the rate of a photochemical reaction.

Mathematically, for a given temperature, a thermal reaction is represented by a constant rate, while the rate of the photochemical transition varies linearly with light intensity. Fig. 13 (a) shows an example of photocycle combining the conclusions from several spectroscopic studies[26, 121] on different channelrhodopsins. The photocycle consists in two separate cycles, each involving a conductive state (P520 and P520′), with possible transitions from one half-cycle to the other and a side reaction (states P380 and P353) with a transition to either of the two half-cycles.

Fig. 13 (b) shows the Markov chain we used to model the behavior of ChrimsonR. It is a simplified version of the more realistic photocycle introduced above. The two half-cycles are represented by $C_1 \rightleftarrows O_1$ and $C_2 \rightleftarrows O_2$. C_1 and C_2 represent the ground states D480 and D470 respectively. The intermediate states P500, P500′, P390 and P390′, which are fast intermediates, have been suppressed. States P520 and P480 have been merged into O_1, P520′ and P480′ into O_2. The side reaction is represented by the single state S, and the transi-

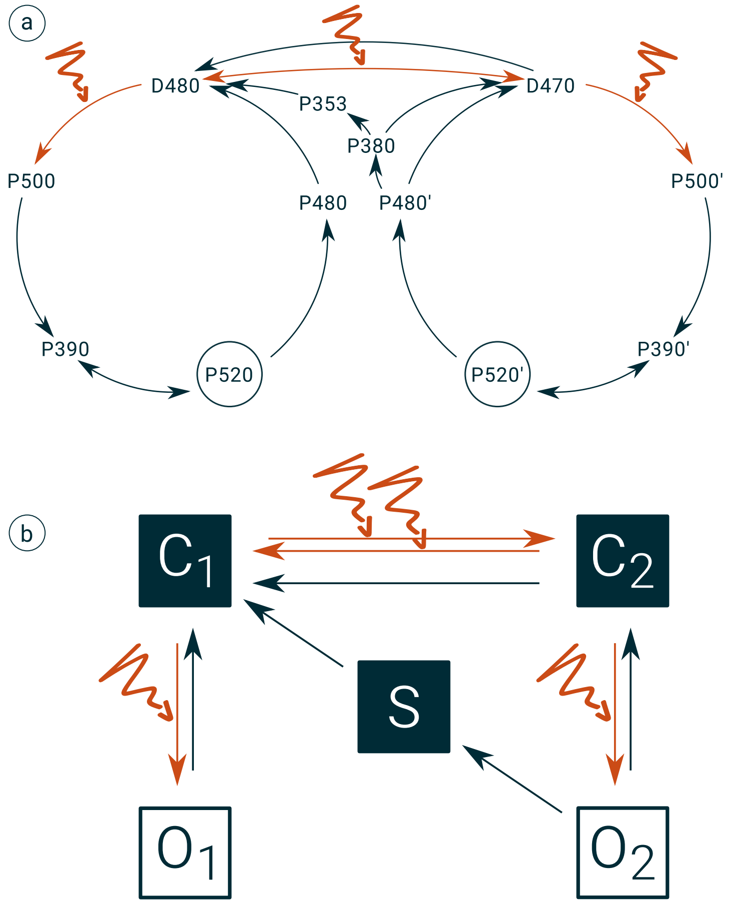

Figure 13: The photocycle (a), such as provided by previous studies on other channelrhodopsins, and the simplified model (b) we used to approximate the electrochemical behavior of ChrimsonR.

tion from S back to half-cycle 2 has been omitted because it could not be resolved based on our data.

Fitting linear combinations of exponential functions A general result on continuous-time Markov chains[42] with time-independent transition rates is that, if the Markov chain has a number N of states, then the probability $p_i(t)$ that the chain is in state $i \in [\![1, N]\!]$ at time t can be expressed as :

$$p_i(t) = \overline{p}_i + \sum_{j=1}^{N-1} \alpha_{i,j} e^{-t/\tau_j} \tag{12}$$

where $(\overline{p}_i)_{1 \leqslant i \leqslant N}$ are the stable probabilities for each state, i.e. the probabilities which are invariant by multiplication by the transition matrix. $(\tau_j)_{1 \leqslant j \leqslant N-1}$ is a set of $N-1$ time constants which are shared by the N states, and $(\alpha_{i,j})_{1 \leqslant i \leqslant N, 1 \leqslant j \leqslant N-1}$ are the weights of the linear combination, and depend on the initial conditions. In our case, the

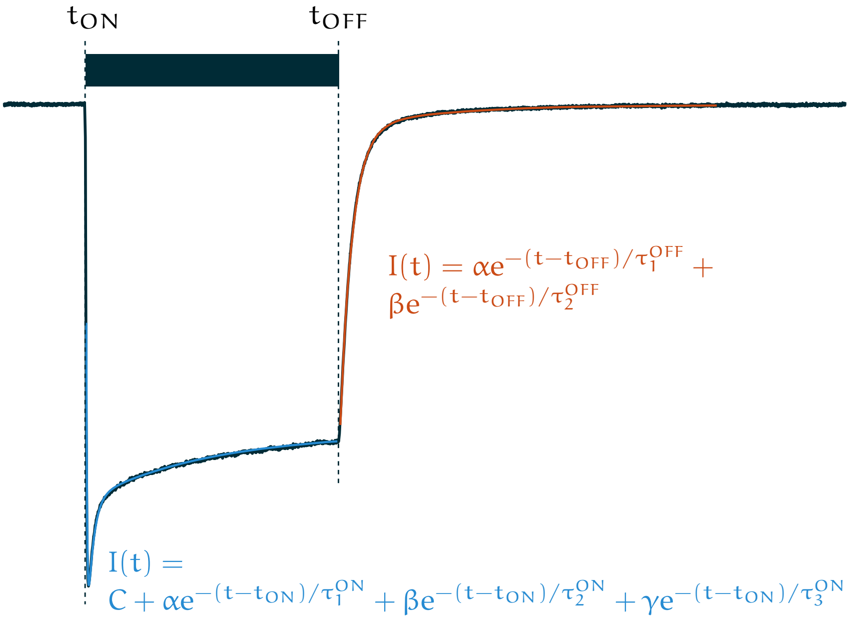

Figure 14: **Estimating the On and Off dynamics** Linear combinations of exponential functions were fitted to the portions of curve where light intensity is constant. The number of terms was fixed based on the shape of the responses. For on dynamics (non-zero light intensity) a constant term is required in order to represent the steady-state level. Given that the steady-state level is known to be zero in the absence of light, the constant term was removed for the analysis of the off dynamics.

transition rates are constant over a time period if and only if light intensity is constant over this period.

Additionally, we do not observe the probabilities directly, but rather a linear combination of the total number of channels which are in each conductive state (P520 and P520′), the weights being the conductances of each state. However, since the number of channels expressed in each patch-clamped cell is constant over time, the fraction of channels in a state i is a good approximation of the probability p_i. Since the N states share the same $N-1$ time constants, the linear combination only affects the coefficients and thus does not prevent us from estimating the time constants.

Fig. 14 shows that three terms describe well the time course of a single pulse during the light pulse, while two terms are enough to describe the dynamics of the decay when light is turned off. This double-exponential dark decay is reported for all channelrhodopsins in the literature. It is in fact the main argument against the single photocycle model, which would be the natural conclusion based on spectroscopic studies only.

Fitting the parameters of the Markov kinetic model The parameters of our model was separated in different groups and fitted sequentially.

First, the two off time constants ($O_1 \rightarrow C_1$ and $O_2 \rightarrow C_2$) were fixed using the numerical results of the analysis on the linear combinations of exponential functions. In this configuration where the two open states are not connected, the values of the two transitions must be the two time constants describing the off dynamics. Analyzes with the two possibilities showed that the fastest time constant must be associated with the stable cycle $C_1 \rightleftarrows O_1$.

Then the other time constants were estimated by minimizing the square difference between the recorded and simulated conductances. The parameters of the error function were the time constants of the transitions, except the transitions $O_1 \rightarrow C_1$ and $O_2 \rightarrow C_2$, fixed in the previous step as well as the thermal transition $C_2 \rightarrow C_1$ which is too slow to be estimated using experiments implemented at this time scale, and which was arbitrarily fixed to a very small value (approx. 30 minutes) so it did not affect the dynamics of the system.

Medium-term and long-term activation curves The side reaction involving the slow state (states P380 and P353 in Fig. 13 (a) or S in Fig. 13 (b)) makes it tricky to estimate the activation curve (steady-state conductance level as a function of the light intensity) of the protein. In fact, it is not trivial to disentangle the influence of the current light intensity level from the history of stimulation on the amplitude of the response. Our estimation of the activation curve relies on the prior fitting procedure of the five-state model to the data. The activation curve can then be recovered through the model, as the steady-state level of the response of the model to the whole range of light intensities.

Here, we present the activation curves at two different time scales. The long-term time scale (several minutes) is enough for the whole chain to converge, including the transitions through the slow state. But we also include a medium-term activation curve, which represents the state of the Markov chain where all transitions but the slow ones have converged. This medium-term equilibrium is obtained by arbitrarily setting the transition rate from O_2 to S to zero.

The analysis of the time constants describing the response of ChrimsonR according to the five-state model allows us to estimate the time constant of the transition from the medium-term to the long-term equilibria.

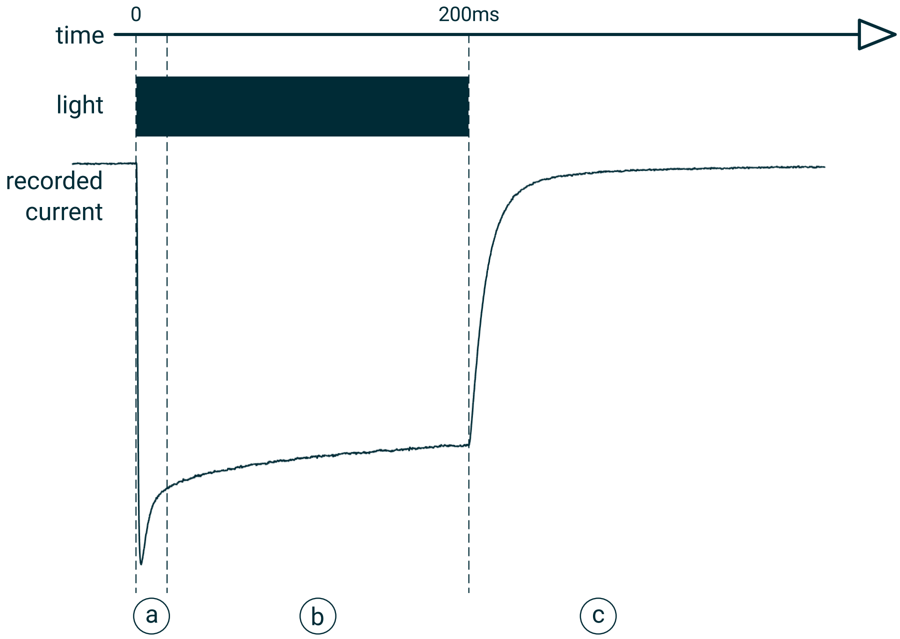

Figure 15: **Generic shape of the response to a square pulse of light**: (a) a fast response, which depends on the stimulation history and lasts from 10 ms to 40 ms, is followed by (b) a slow decay towards a non-zero equilibrium value. (c) When light is turned off, the recorded current converges towards zero with a double-exponential decay.

3.2 RESULTS

Main observations

In response to a single 200 ms light pulse, the typical response that we observe has three main features : (i) a fast response occurring in the first 10 ms to 40 ms from pulse onset which can either be an overshoot or an undershoot depending on several factors detailed below, followed by (ii) a slow decay with a time constant of the order of 100 ms (see paragraph *On kinetics*) towards a non-zero equilibrium value which is not reached within the 200 ms, and (iii) when light is switched off, a decay towards zero characterized by the sum of two exponential terms.

Additionally, when the pulse is repeated (every 2 s in our protocols) at high light intensity ($I \geqslant 10^{17}\,\mathrm{ph\,s^{-1}\,cm^{-2}}$) the response for a given pulse is generally of the same shape as the previous one but with a slightly smaller amplitude.

Our most valuable observations are derived from clean recordings of the same cell over several minutes, during which the building blocks of the protocols (series of ten pulses) were not arranged in a monotonic order, i.e. block intensity goes up and down (Fig. 22, Fig. 24 and Fig. 25).

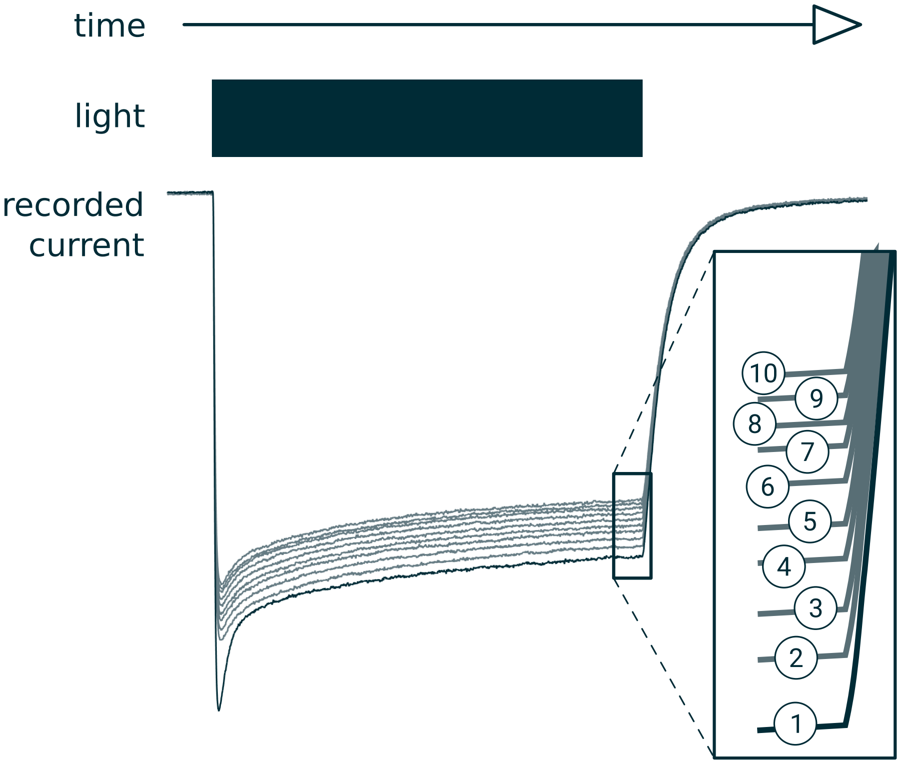

Figure 16: **Responses to repeated stimuli** (200 ms light pulses at a fixed intensity — $I = 1.17 \times 10^{19}\,\mathrm{ph\,s^{-1}\,cm^{-2}}$ — one pulse every 2 s) : the ten responses are synchronized with respect to the stimulus and stacked. The order of the sweeps is given in the box on the right. Except the first pulse, the nine others are only scaled versions of each other.

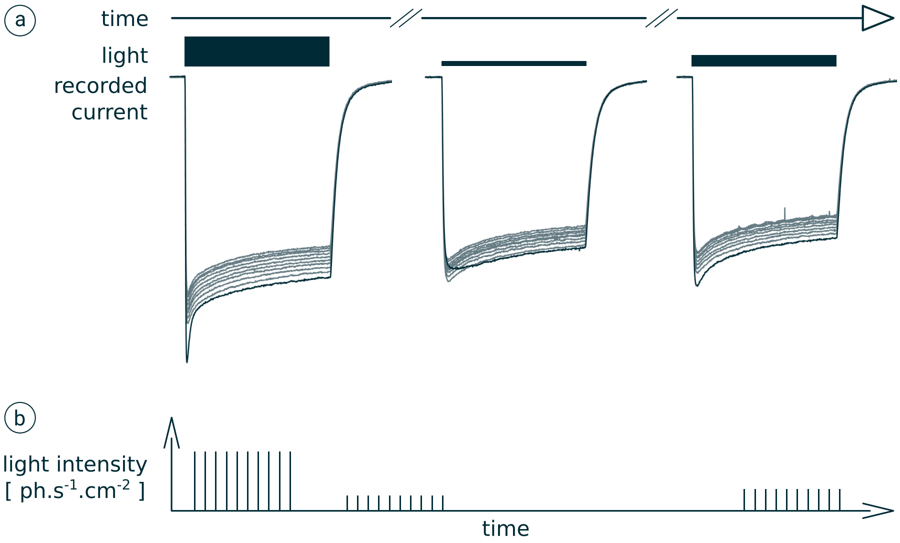

Figure 17: **Responses to series of ten** 200 ms **pulses**: (a) Stacked responses : the response to the first pulse is darker. (b) Stimulus description : it consists in three series of ten pulses. In each series, light intensity is constant and one pulse is emitted every 2 s. Light intensities are respectively 1.17×10^{19}, 3.12×10^{18} and 4.36×10^{18} ph s^{-1} cm^{-2}.

In this situation, the shape of the fast part of the response (occurring in the first 10 ms to 40 ms) depends substantially on the light intensity of the previous pulse. Specifically, for a give pulse with a given light intensity, the higher the light intensity during the previous pulse, the lower the fast part of the response (see Fig. 17 (a)). It is our understanding, that the time interval between the two pulses has little influence on the fast part of the response. In contrast, the slow part of the response is unaffected by the previous stimulation patterns.

Time constants

On kinetics At low light levels ($I \leqslant 5 \times 10^{17}$ ph s^{-1} cm^{-2}), the time course of the recorded current during an interval when light is on and of constant intensity is well approximated by the linear combination of (i) a constant term representing the limit current and (ii) two exponential terms with distinct time constants. At higher light levels, an additional exponential term, with a third time constant, is necessary to describe accurately the response to a single light pulse.

Fig. 18 (b) shows the estimated time constants as a function of light intensity. The fastest time constant (often referred to as τ_{ON} in the literature, black curve) — which represents the time scale at which the protein can be activated — is well estimated. It decreases with light intensity, i.e. the higher the intensity, the faster the protein is activated. The estimation of the second time constant is less precise (Fig. 18 (b), blue curve). The estimation of the third time constant (orange curve) is very erratic, as shown by the very large confidence intervals on most of the recoded cells on Fig. 18 (a). However, the analyzes run on the cleanest data sets suggest that a time constant between 10^2 ms and 2×10^2 ms, relatively independent of the stimulation intensity is a good initial guess.

Off kinetics In agreement with what was described in the literature, our data shows that the dynamics of the observed current when light is turned off is very well approximated by a linear combination of two exponential terms. Fig. 19 shows the numerical results obtained separately on each recorded cell.

Slow adaptation process

As stated in the *Main observations* section, our data shows that the behavior of ChrimsonR involves a slow adaptation process such that the overall amplitude of the response to a series of identical pulses decreases with each pulse. This feature has not been previously reported in electrophysiological experiments, but we think that this phenomenon can be compared with the side reaction involving states P380 and P353 of the photocycle of the channelrhodopsin C128T mutant described in particular in [130, 121].

The time scale at which this adaptation occurs is too long to be observed on single 200 ms pulses. Therefore, the analysis of the on-dynamics (paragraph *On kinetics*) cannot bring any valuable information to estimate the time constant of this phenomenon.

Instead, we used the Markov kinetic model involving a slow state representing the side reaction and we fitted the model on the responses of the same cell to stimulation protocols consisting of repeated series of ten pulses. Since the protocols span over several minutes, they allow the estimation of this adaptation process, especially the recovery time constant. Fig. 27 shows side by side the time constants such as estimated directly on single pulses and theoretically using the model after parameter estimation. This second method reveals a fourth time constant, much longer than the three others, and characterizing the slow adaptation process.

Simplified photocycles

Global approximation The model aims at capturing parts (b) and (c) of the typical response shown in Fig. 15, i.e. the overall amplitude of the response including the slow adaption with a time constant of the order of 10^2 ms and the double exponential decay.

As stated in Section 3.1, some parameters were fitted individually for each cell, like the time constants from the open states back to

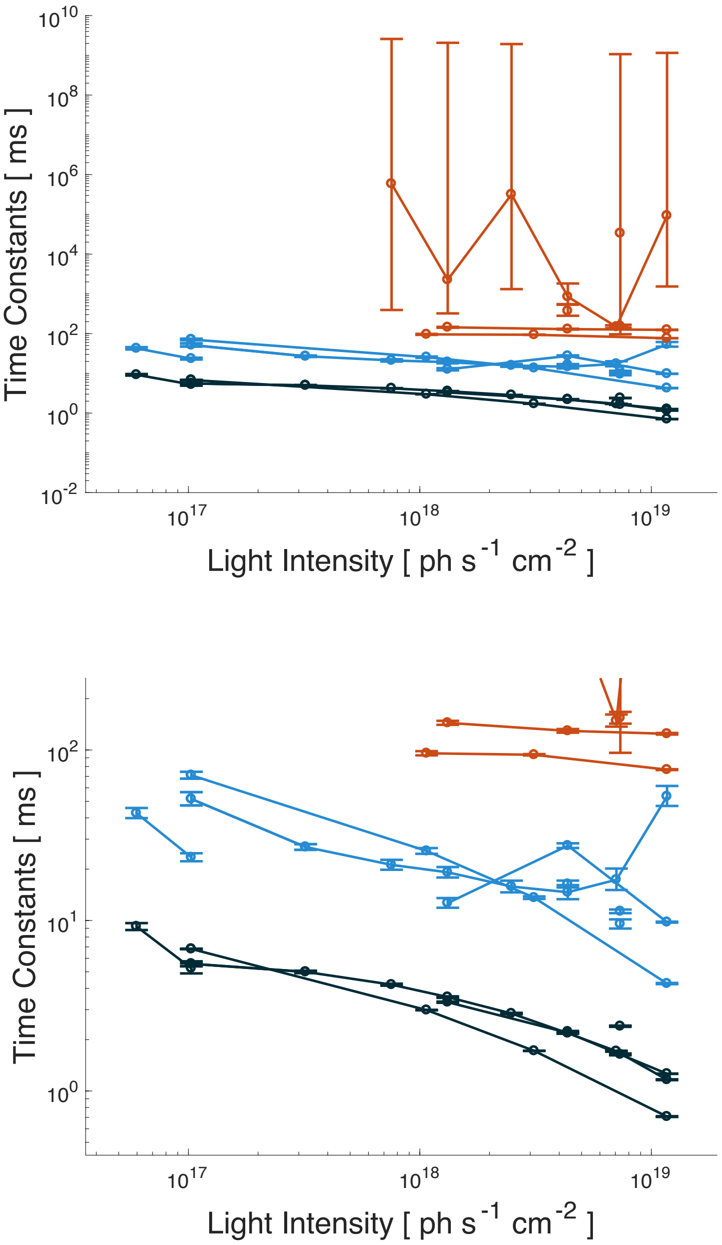

Figure 18: **On kinetics analysis of experimental data**: linear combinations of exponential functions have been fitted on the recorded responses to light pulses. Top : when available, all three time constants are shown. The precision on the slowest time constant is usually very poor. Bottom : same curve displaying only the two fastest time constants. The circles show the estimated values for each time constant, and the error bar show the 95% confidence intervals. The data points that are linked together belong to the same cell.

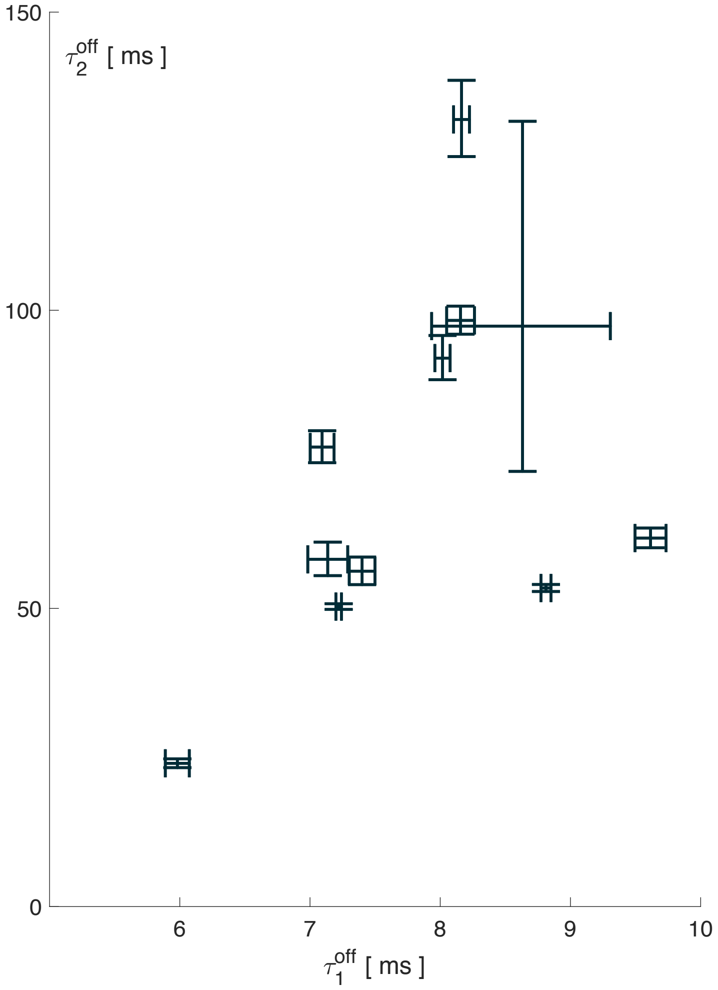

Figure 19: **Numerical results for the estimation of the time constants governing the off dynamics of ChrimsonR-tdTomato**. For each cell, the estimations of the fast (abscissa) and slow (ordinate) time constants are shown along with the 95% confidence intervals. It is worth noting that intra-cell variability is generally much lower than the inter-cell variability.

Table 1: Estimated conductances for the six different cells (in nS)

Cell	g_1	g_2
1	17	3.3
2	3.92	0.44
3	4.85	2.01
4	3.19	0.39
5	15.4	2.88
6	8.07	2.84

Table 2: Estimated photochemical transition rates (in ms^{-1} / $\mathrm{ph\,s}^{-1}\,\mathrm{cm}^{-2}$)

Cell	$C_1 \rightarrow O_1$	$C_2 \rightarrow O_2$	$C_1 \rightarrow C_2$	$C_2 \rightarrow C_1$
1	1.61×10^{-19}	3.03×10^{-20}	1.16×10^{-20}	5.99×10^{-20}
2	1.67×10^{-19}	3.89×10^{-20}	1.28×10^{-20}	6.17×10^{-20}
3	1.15×10^{-19}	5.84×10^{-20}	2.96×10^{-21}	4.07×10^{-20}
4	4.60×10^{-19}	1.23×10^{-19}	5.13×10^{-20}	1.46×10^{-19}
5	1.10×10^{-19}	7.20×10^{-20}	1.94×10^{-21}	1.44×10^{-20}
6	1.60×10^{-19}	3.87×10^{-20}	4.93×10^{-20}	1.75×10^{-19}

the closed states, other parameters were fitted on the most relevant data sets and then used for all parameter sets. Numerical values are provided in Table 1, Table 2 and Table 3.

Fig. 20 to Fig. 25 display the comparison between the data and the simulations for six different cells. They show that the main features are in fact captured by the five-state model.

Activation curve The five-state model allows the recovery of important information which in not directly available through simple analysis of the raw data. One such information is the activation curve of the protein, i.e. the amplitude of the steady-state conductance of the protein as a function of light intensity.

Table 3: Estimated thermal transition rates (in ms^{-1})

Cell	$O_1 \rightarrow C_1$	$O_2 \rightarrow C_2$	$O_2 \rightarrow S$	$C_2 \rightarrow C_1$	$S \rightarrow C_1$
1	0.14	1.14×10^{-2}	48.1×10^{-5}	10^{-7}	5.91×10^{-6}
2	0.12	1.78×10^{-2}	7.89×10^{-5}	10^{-7}	3×10^{-6}
3	0.13	1.78×10^{-2}	8.56×10^{-5}	10^{-7}	3×10^{-6}
4	0.10	1.40×10^{-2}	11.5×10^{-5}	10^{-7}	2.49×10^{-6}
5	0.12	1.78×10^{-2}	8.79×10^{-5}	10^{-7}	3×10^{-6}
6	0.10	1.37×10^{-2}	19.2×10^{-5}	10^{-7}	1.5×10^{-6}

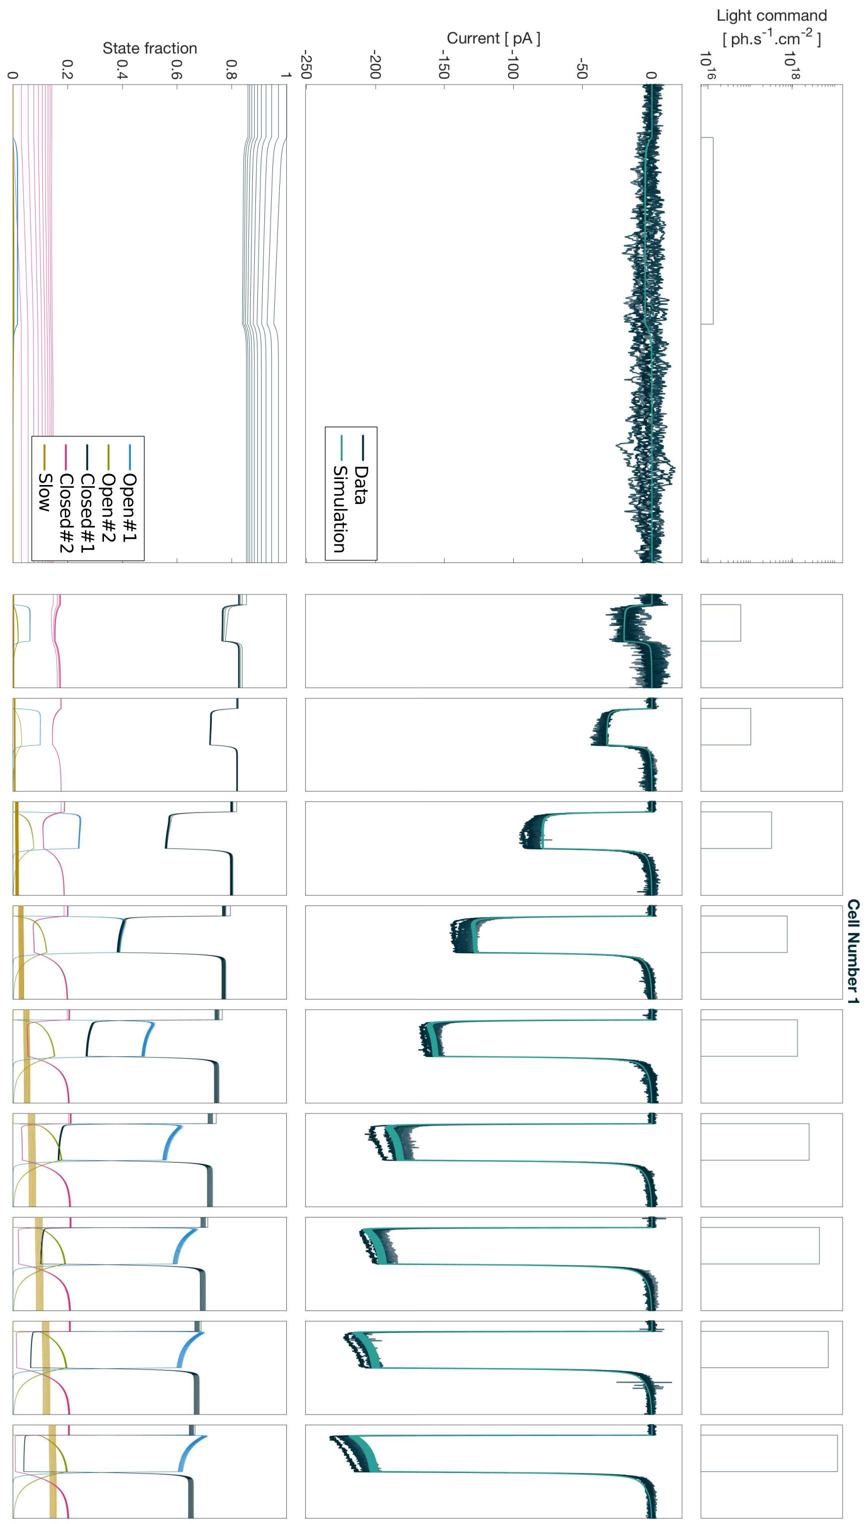

Figure 20: **Data and simulation comparison Cell 1:** Each column represents the repetition of ten superimposed 200 ms pulses with the same intensity. The first sequence of ten pulses is shown in a larger window to increase visibility. The top row represents the stimulus, in log scale. The second row shows the recorded (dark) and simulated (cyan) currents. The ten responses are superimposed and synchronized with respect to the stimulation pulse. The bottom row shows the detail of the evolution of the fraction of channels in each of the five states of the Markov kinetic model.

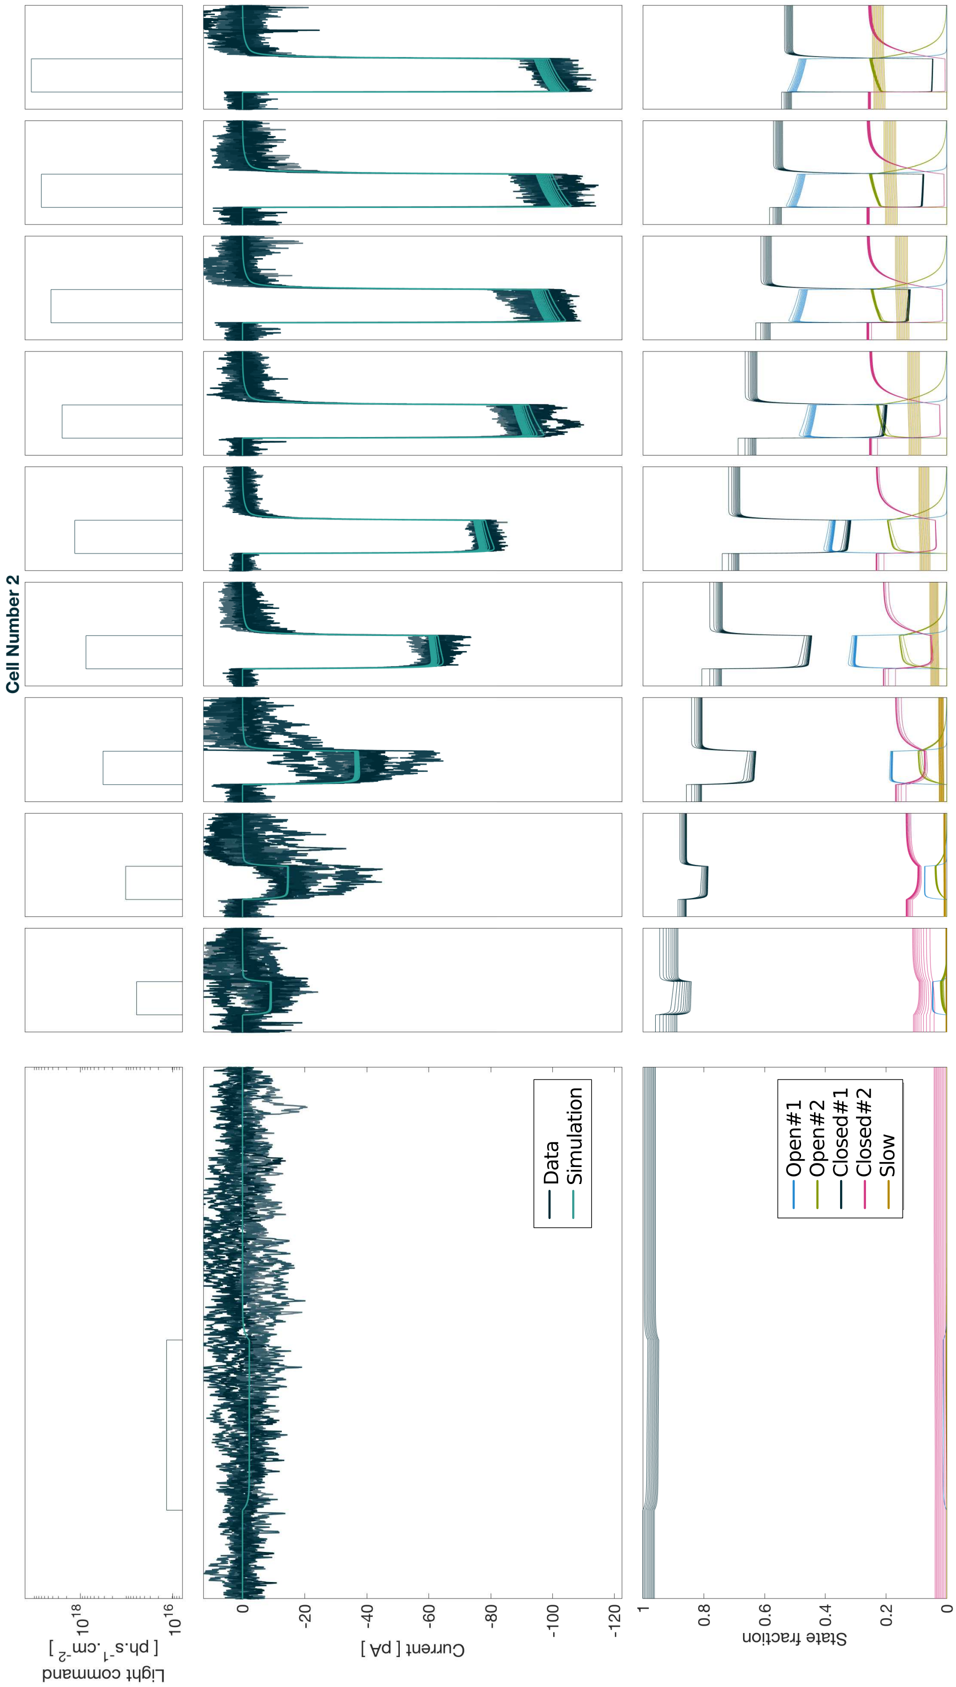

Figure 21: **Data and simulation comparison** Cell 2

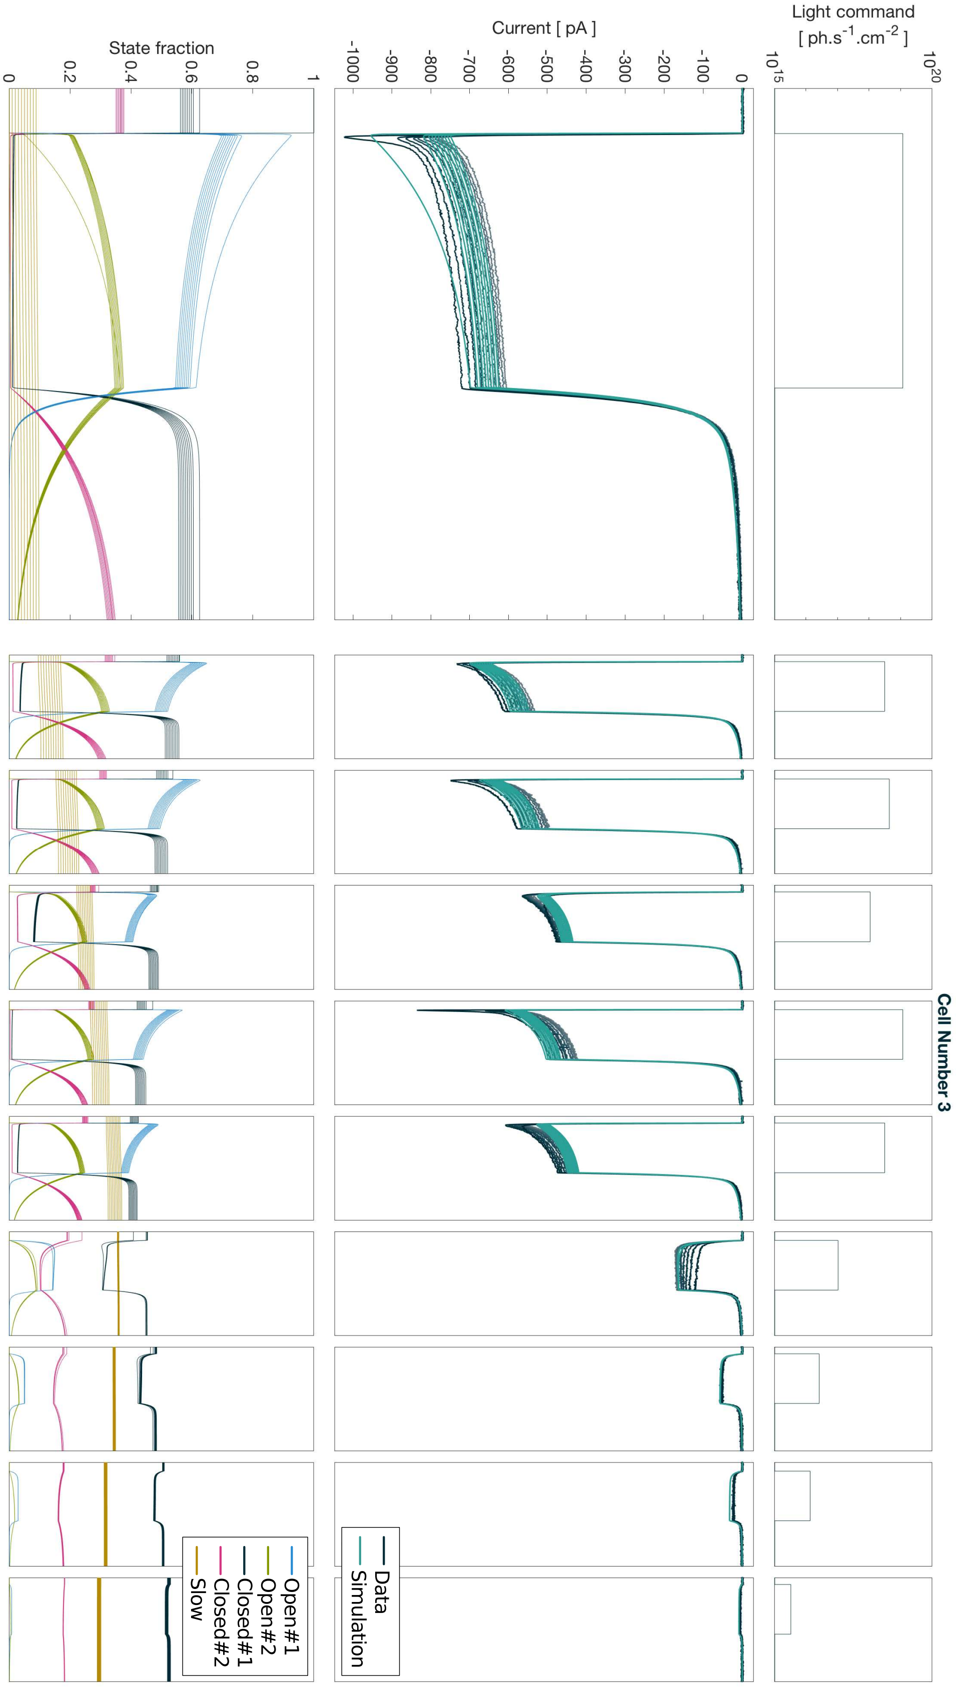

Figure 22: **Data and simulation comparison** Cell 3

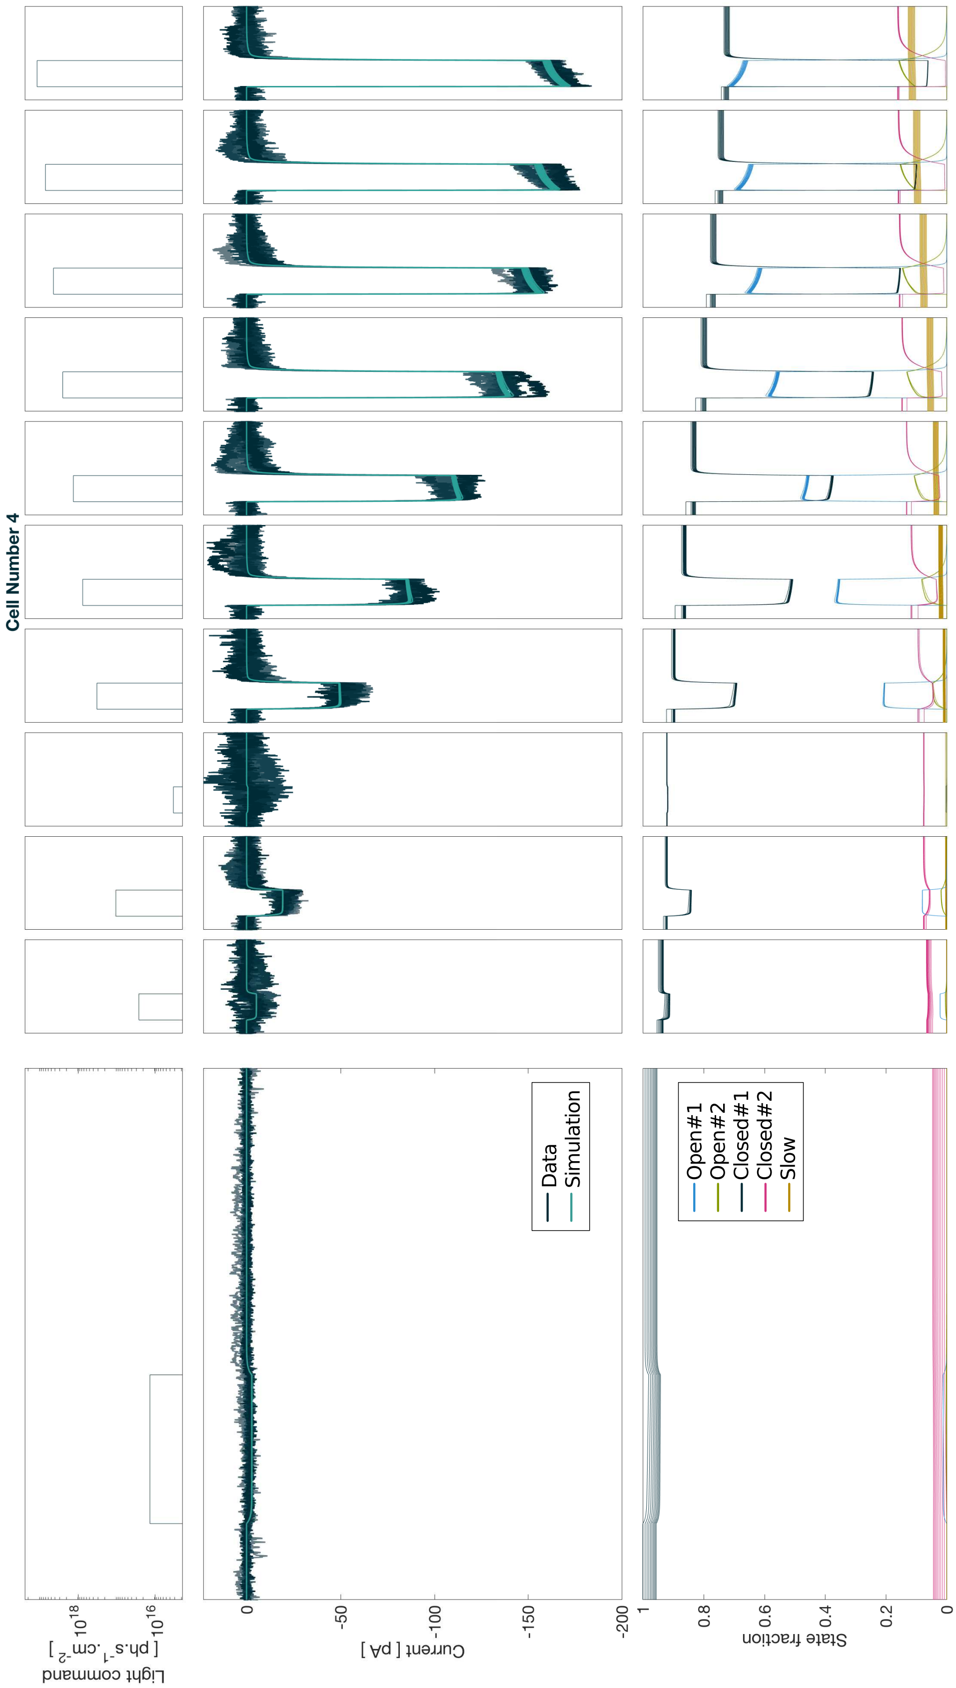

Figure 23: **Data and simulation comparison** Cell 4

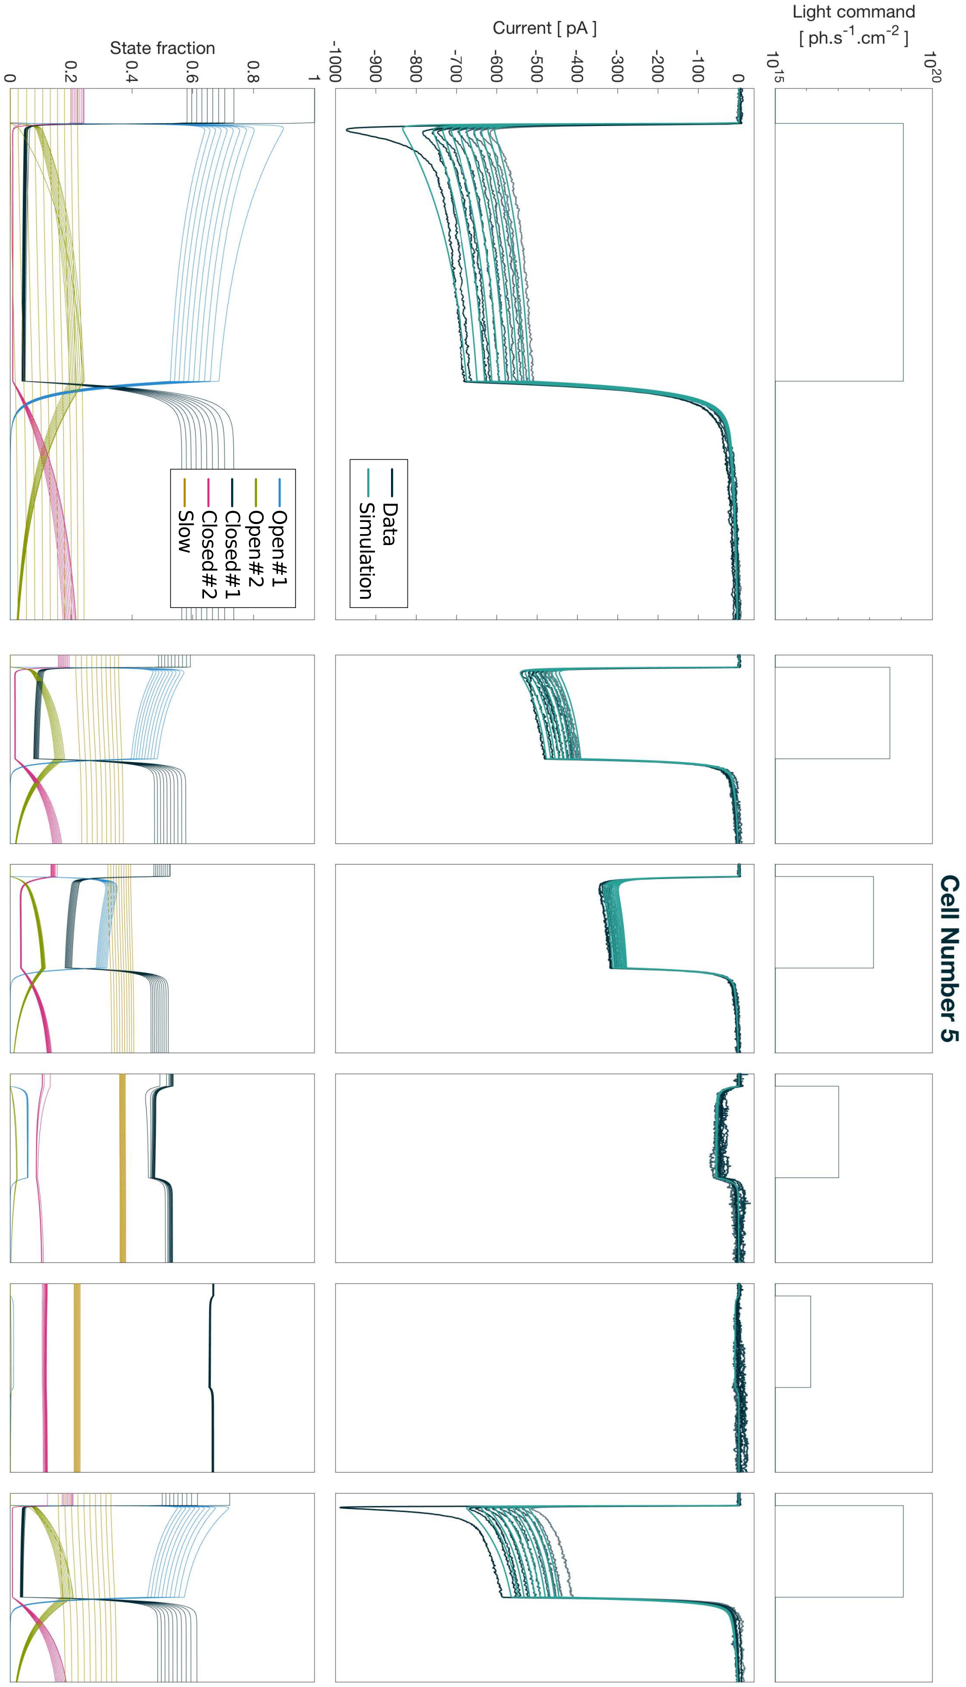

Figure 24: **Data and simulation comparison** Cell 5

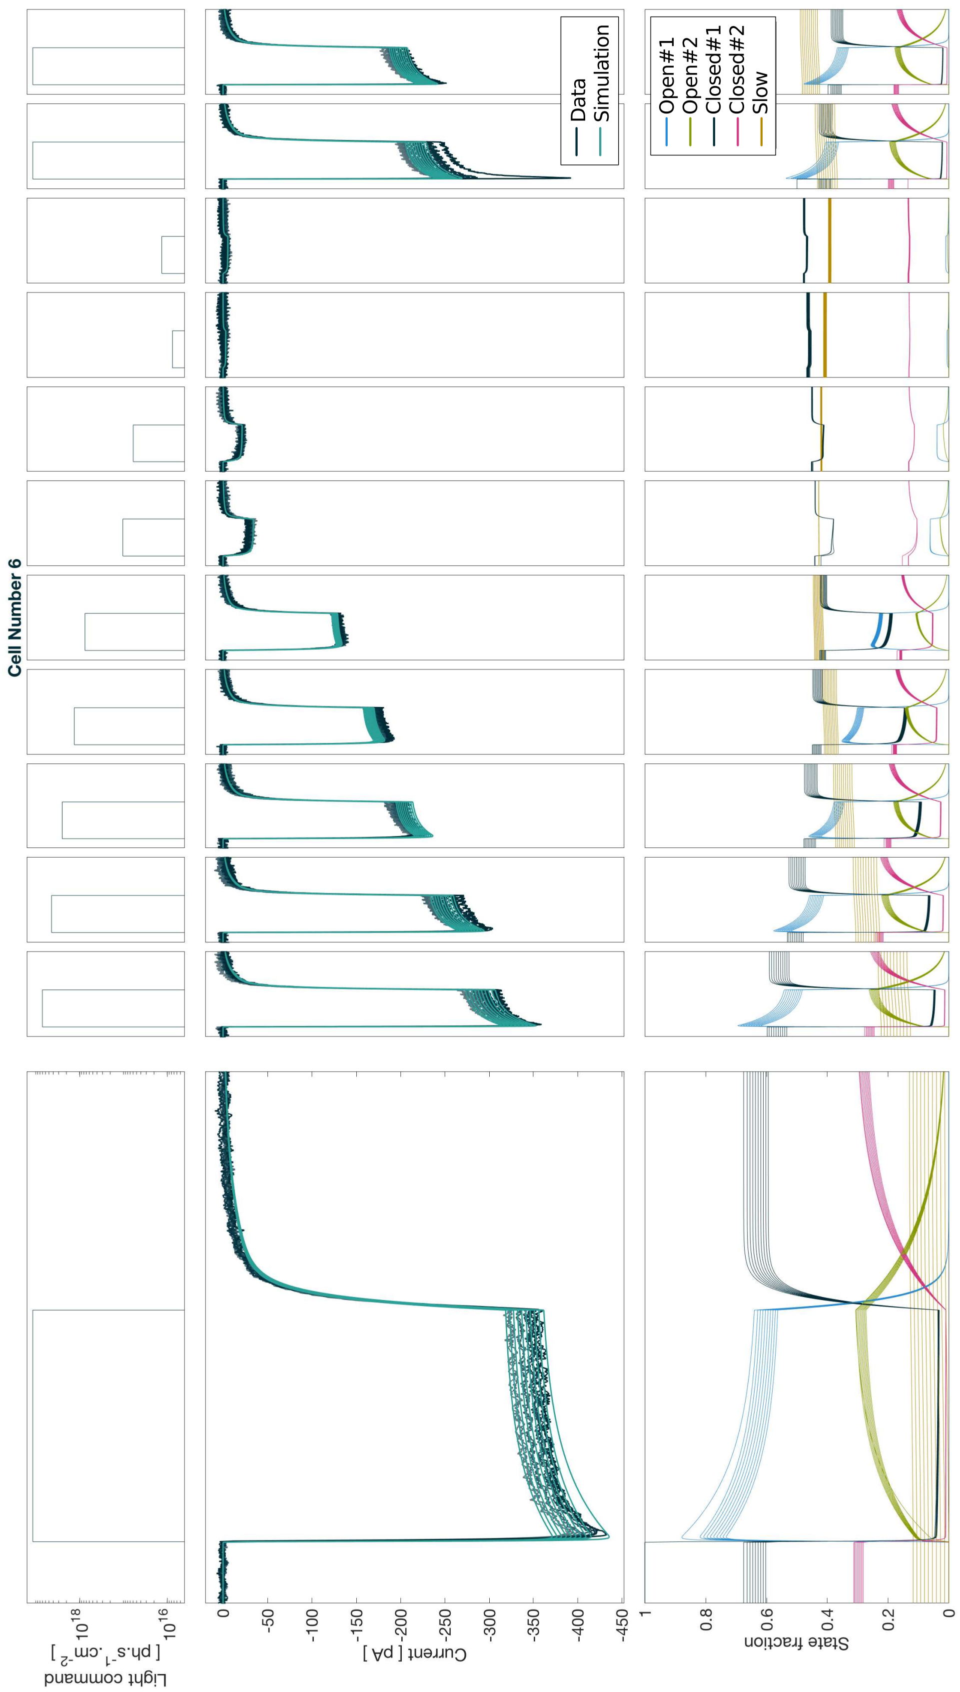

Figure 25: **Data and simulation comparison** Cell 6

Fig. 26 shows the activation curves for all parameters sets presented above. For each parameter set, two distinct activation curves are presented, the first one is the steady-state response at an intermediate time scale around 1 s for which the conductance has apparently reached a plateau, but which in fact slowly decays to a lower plateau due to the side reaction represented by the state S in Fig. 13 (b). The second activation curve represents this real steady-state, including the side reaction.

The estimation of the slow time constant of the transition from the slow state S to the stable closed state C_1 is quite imprecise because it is long relative to the time of the experiment and that other phenomena involved in the experiment — clamping condition, physiological state of the cell, etc. . . — occur at the same time scale. However, assuming a sensible numerical value of 5 minutes for this time constant, we estimate that the long-term steady-state response is somewhere between 3 and 15% of the medium-term value, the one that is observed when proceeding to a short experiment.

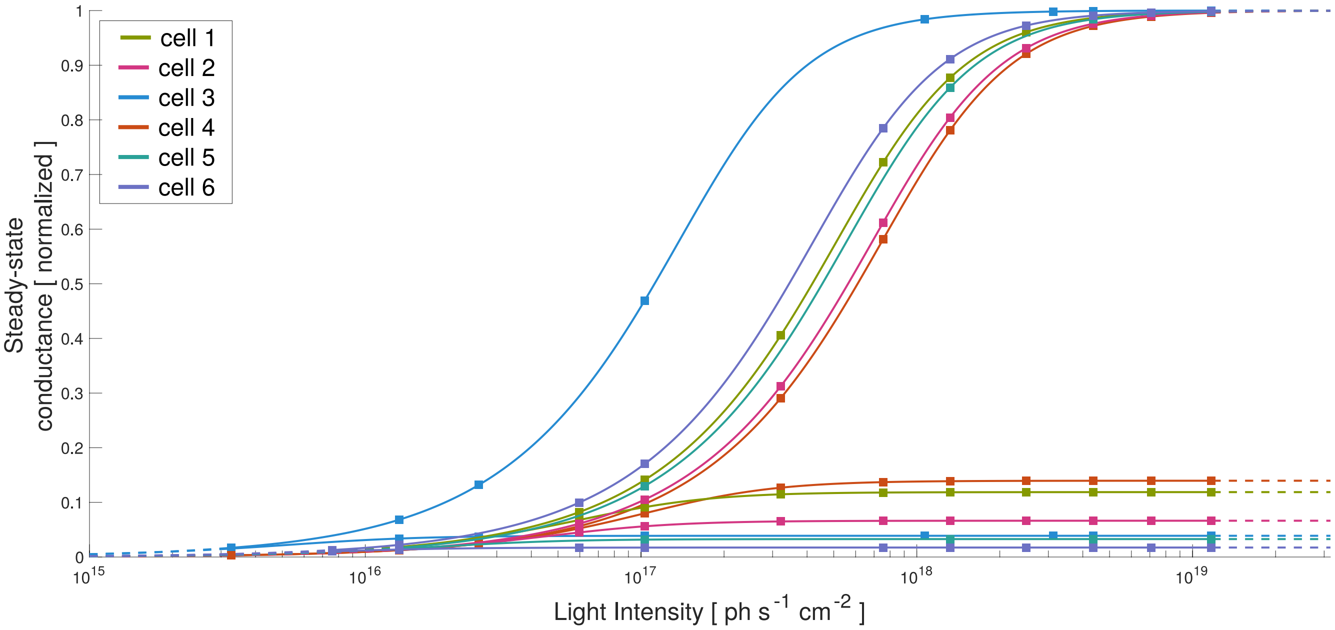

Figure 26: **Activation curve of ChrimsonR at medium- and long-term steady-state** The curves are derived from the Five-state model fitted on data from four different cells, for which the response was observed on a wide range of light intensities. The curves were obtained theoretically according to the procedure described in the *Activation curve* paragraph and are normalized with respect to the limit value of medium term activation curve. The dots on each curve represent the light intensity values which were present in the stimulation protocol on which the model was fitted. Thus, the curve is shown in dotted line outside the range of actual stimulation values. ChrimsonR starts to respond significantly around $10^{16}\,\mathrm{ph\,s^{-1}\,cm^{-2}}$ and saturates above $10^{18}\,\mathrm{ph\,s^{-1}\,cm^{-2}}$. Beyond this limit, the steady-state value of the conductance is unaffected. Only the dynamics are impacted.

ON time constants comparison Contrary to the off dynamics, the preliminary studies of the on dynamics of the responses were not used to fit the parameters of the five-state model. As a safety check, we display side by side in Fig. 27 (i) the time constants estimated directly on the raw data as presented in paragraph *On kinetics* and (ii) the time constants of the five sets of parameters introduced above.

The figure shows that the dynamics are comparable between the recorded and simulated responses on the range of recorded light intensities (black, blue and orange curves). The time constants derived directly from the model are slightly quicker than the ones estimated on the data. Computing the time constants through the model allows us to reveal the dynamics at with the open states equilibrate with the slow state S (green curve). It represents the time scale at which we slide from the medium-term to the long-term steady-state (see Fig. 26).

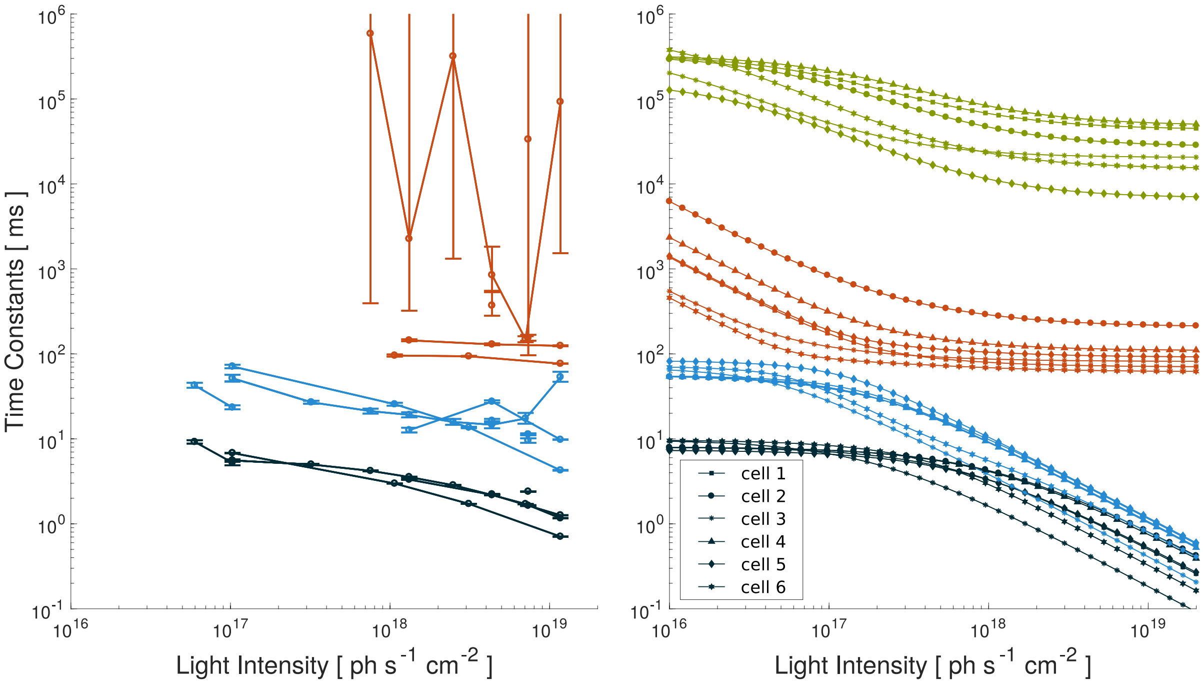

Figure 27: **On dynamics comparison** Time constants of the linear combination of exponential functions underlying the responses to square pulses of light stimulation. (left) time constants are estimated directly on the data. Depending on the intensity of the light command, either two or three distinct time constants are estimated. (right) time constants derived from the five-state model dynamics. The time constants are computed as the non-zero eigen values of the differential operator governing the evolution of the probability distribution of the Markov chain. Note that this method reveals a fourth time constant. This fourth time constant represents the dynamics at which the open states equilibrate with the slow state S.

3.3 DISCUSSION

In this paper, we propose a five-state Markov kinetic model describing the conductance response of channelrhodopsins ChrimsonR to an arbitrary light stimulus. The model was chosen to capture the main features of the experimental observations : activation curve, on and off kinetics, medium term and long term light adaptation. The model was designed in the direct lineage of similar work, carried out over the past decade[63, 105, 58, 47], but also taking into account the more recent conclusion on the actual, more complex, photocycle of channelrhodopsins[26].

Our analysis is based on a set of voltage-clamp experiments on ChrimsonR-transfected HEK293 cells at a holding potential of $-60\,\mathrm{mV}$ and at physiological temperature for direct use in medical application. Direct observation of the data revealed that there is a fast light adaptation mechanism that has not been described previously. It was already known, from the shape of the conductance responses to light pulses, that the conductance does not reach a plateau directly, but rather peaks first before decaying to a stable value. This decay has been described as light adaptation and corresponds to an equilibration between the two cycles. We additionally report that the intensity of the light pulse influences the ratio between the two cycles, and in turn the resulting ratio between the populations in the two closed states D470 and D480 after light is turned off. Finally this ratio determines the shape of the first part of the response to a new light pulse.

However, this feature is not captured by our five-state model, and we did not manage to capture it with any of the numerous n-state models that we tested. This underlines that the complex photocycle of channelrhodopsin is not fully understood and that work remains to be done before we establish the complete outline of the photocycle. However, we believe that once the photocycle is understood, fitting it to voltage-clamp data of the different known channelrhodopsin is going to be quite straightforward. We also believe that simulations of voltage-clamp data is a major tool for unveiling the photocycle, along with spectroscopy (which provides valuable information regarding the time constants of the existing reactions) and crystallography (which identifies the possible chemical reactions, changes of conformation and associate them with the structure of the protein).

The model that we propose, along with the parameters sets fitted on different cells, captures the main features of the electrophysiological properties of ChrimsonR. It is a useful tool for predicting the conductance of the protein in response to an arbitrary light stimulation, in neurophysiological experiments or in medical applications such as retinal prostheses. Since the model was built on experiments conducted at physiological temperature, the model is particularly useful for medical applications and is directly applicable.

The main limitation of our model is that is does not capture the fast light adaptation mechanism. For that reason, its prediction power will decrease as the temporal frequencies of the stimulation exceeds approximately 100 Hz. However, the model displays the correct ON dynamics on the whole range of relevant stimulation intensities, therefore the prediction will still be close to the actual signal even for high frequencies. Other limitations come from the fact that we have not studied ChrimsonR's spectral sensitivity and the ion selectivity of the two open states of the protein.

4

MODELING THE SPIKING BEHAVIOR OF THE TRANSFECTED CELL

4.1 INTRODUCTION

The retina transmits information to the higher areas of the visual system via the parallel spike trains propagated along the optic nerve. In the case of an RP patient treated in the context of the GS030 project, these spike trains travelling along the optic nerve are the result of the light stimulation of the ChrimsonR-expressing Retinal Ganglion Cells. The downstream visual pathways will perform the same operations they used to carry out before the degeneration of the retina on the artificially induced spike trains, and the output of these operations will dictate what the patient perceives.

In this chapter, we develop a mathematical model of the spike generation process, the mechanism that transforms the stimulation signal (the value for the light intensity varying with the time) into a spike train. The goal of this model is to serve as a tool to analyze quantitatively the transfer of information between (i) the light pattern which can be controlled and (ii) the spike trains which are the support for the visual information that the downstream visual areas of the Central Nervous System (CNS) receive. In other words, this work models the behavior of the Brain-Machine Interface so that the preprocessing steps of the device information pipeline can be optimized in order to maximize the information rate that is actually transmitted through the interface.

Additionally, preliminary analyses in primates suggest that Retinal Ganglion Cells in the perifoveal region are the primary targets of transfection. According to a study by Dacey[37] on 46 human retinas, an "estimate of total midget cell density suggested that the proportion of midget cells increased from about 45% of total ganglion cell density in the retinal periphery to about 95% in the central retina". Therefore, the study presented in this chapter is limited to this type of RGC. Midget Ganglion Cells come in two subtypes, ON and OFF. These two types differ in many ways: they do not stratify in the same portions of the inner plexiform layer and do not have the same dendritic field sizes. More importantly, ON and OFF Midget Ganglion Cells code for the opposing local features of the visual scene. However, due to their similar morphologies[37], we will make the assumption that both will respond similarly to photostimulation.

A central question to this study is the shape of the model that we should use to predict the behavior of the transfected cells. Other studies on the spiking behavior of neurons in the early visual system, such as the one by Keat et al. stressed that neuronal responses can be very precise, and that a given stimulation pattern would evoke very similar spike trains if repeated, in which individual action potentials would be reliably triggered with a millisecond temporal precision : "neurons in the early visual system — from the retina to the lateral geniculate nucleus to area V1 — can deliver remarkably reproducible spike trains, whose trial-to-trial variability is clearly lower than predicted from the simple firing rate formalism. [. . .] Consequently, one needs a different framework for studying this neural code that allows for the prediction of individual spikes and spike patterns with high timing accuracy, but also accounts for the remaining stochastic variability in these responses."

This temporal resolution of the order of the millisecond has also been observed in mammal's early downstream visual areas. such as the LGN[138, 120, 90, 82] and the Middle Temporal (MT) area of the visual cortex[27]. Additionally, [30] showed that the fine temporal resolution of the spike train had a functional role in encoding the visual stimulus, even if the latter varies with slower dynamics.

This argues in favor of building a model able to capture the precise structure of the spike trains, i.e. the timing of each individual spike, rather than the spike rate smoothed over a longer timescale.

Our choice fell on a variant of the Linear-NonLinear Integrate-and-Fire (L-LN-IF) model introduced by Pillow et al.. This model offers several benefits which are relevant to our approach. First, as demonstrated by the original paper, the model is able to reproduce the fine temporal organization of the recorded spike trains. This is the primary requirement for the work we pursue in this study. The second interesting feature is the modeling of a noise term. When dealing with information transmission and encoding, noise is always a key parameter since it dictates the boundary on the precision up to which information can encoded. Trying to encode more information can only result in ambiguous messages.

Since the primary purpose of our model is to provide an analytical tool to study information transmission through the BCI, we decided to settle for a simple functional model such as the L-LN-IF model rather than a more complex — and also more realistic — conductance-based model. This choice is substantiated by the study by Jolivet et al.. Their paper shows that *Generalized Integrate-and-Fire models* can *de facto* approximate the spike trains of more detailed (conductance-based) models with a high accuracy. However, they point out that the goodness of the approximation is usually restricted to a class of stimulation patterns, which obviously imposes caution on the future use of the model.

We made a few modifications to the original model presented in [113]. First, we have the advantage to know precisely how light interacts with the cell. In fact, spikes originate from the graded depolarizations induced by the opening of the light-gated ion channels ChrimsonR. In this chapter, we make the approximation that the membrane potential is directly driven by the conductance of the population of ChrimsonR proteins. Instead of modeling precisely the flow of ions through the channels as a function of the membrane potential and their Nernst equilibrium potential, we make the simplifying assumption that a constant conductance of the channel population leads to a constant current through the membrane. This is equivalent to omitting the impact of the fluctuations of the subthreshold membrane potential on the force driving the ions though the open channels. There are several reasons for this choice, including (i) our lack of knowledge about the selectivity of the channelrhodopsin ChrimsonR and (ii) the rest of the model is below this level of detail.

The second simplification is to choose a simpler feedback kernel basis. In [113], the authors use a stimulus kernel and a spike feedback kernel in order to capture how the RGC encodes complex spatiotemporal features of the stimulus. In our situation, the upstream circuitry of the retina is not involved in the spike generation process, leading to a much simpler transformation between the light stimulus and the spike trains (once we know how the ChrimsonR population responds).

4.2 METHODS

Animals

All experiments were done in accordance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals. The protocol was approved by the local animal ethics committees and was conducted in accordance with Directive 2010/63/EU of the European Parliament. All animals used in this study were cynomolgus macaques (Macaca fascicularis) of foreign origin.

AAV Production

SNCG promoter was cloned into an AAV backbone plasmid containing the ChrimsonR sequence in fusion with GFP and td-Tomato. The constructs all included woodchuck hepatitis virus post-transcriptional regulatory element (WPRE) and bovine growth hormone poly(A). Recombinant AAVs were produced by the plasmid co-transfection method[56], and the resulting lysates were purified via iodixanol gradient ultracentrifugation as previously described. Briefly, a 40% iodixanol fraction was concentrated and buffer was exchanged using Amicon Ultra-15 centrifugal filter units. Vector stocks were then titered for DNase-resistant vector genomes by real-time PCR relative to a standard[8].

Injections

The primates were anesthetized with 10:1 $\mathrm{mg\,kg^{-1}}$ ketamine/xylazine. $100\,\mathrm{mL}$ of viral vector containing either 10^{11}, 5×10^{11} or 10^{12} viral particles was injected into the vitreous. An ophthalmic steroid and antibiotic ointment was applied to the cornea post-injection.

Two-Photon Imaging and Electrophysiological Recordings

A custom-made two-photon microscope equipped with a 25 water immersion objective (XLPLN25xWMP/NA1.05; Olympus) with a pulsed femtosecond laser (InSight DeepSee; Newport Corporation) was used for imaging ChrimsonR-GFP-tdTomdato positive retinal ganglion cells. For live two-photon imaging, retinas were placed in the recording chamber of the microscope, and z stacks were acquired using the excitation laser at a wavelength of 930 nm. Images were processed offline using ImageJ software (NIH). During imaging, the retina was superfused with oxygenized Ames' medium (Sigma-Aldrich).

We used an Axon Multiclamp 700B amplifier for whole-cell patch-clamp and cell-attached recordings. Patch-clamp electrodes were made from borosilicate glass (BF100-50-10; Sutter Instruments) and pulled to 8-10 $\mathrm{M\Omega}$. Pipettes were filled with 112.5 mM $CsMeSO_4$, 1 mM $MgSO_4$, 7.8×10^{-3} mM $CaCl_2$, 0.5 mM 1,2-bis(O-aminophe- noxy)ethane-N, N,N′,N′-tetraacetic acid (BAPTA), 10 mM HEPES, 4 mM ATP-Na_2, 0.5 mM GTP-Na_3, and 5 mM lidocaine N-ethyl bromide (QX314-Br) (pH 7.2). Cells were clamped at a potential of 60 mV to isolate excitatory currents. Cell-attached recordings were obtained in current-clamp configuration (current zero) with electrodes filled with Ames' solution. L-AP4 was added to Ames' medium during all electrophysiological recordings. Retinas were dark-adapted for one hour prior to recordings.

Photostimulation

Photostimulation was performed using a monochromatic light source (Polychrome V, TILL photonics (FEI), Hillsboro, OR). Stimulation was performed at two intensities 1.6×10^{16} and 3.15×10^{17} $\mathrm{ph\,s^{-1}\,cm^{-2}}$. The stimulation was always uniform over the whole recorded cell, and varied with time according to pattern described below. Stimuli were generated using custom-written software in LabVIEW (National Instruments, Austin, TX). Output light intensities were calibrated (10^{12} - 10^{18} $\mathrm{ph\,s^{-1}\,cm^{-2}}$) using a spectrophotometer (USB2000+, Ocean Optics, Dunedin, FL).

Comparison with HEK293 voltage-clamp data

As stated in the introduction, we use our ChrimsonR model from Chapter 3 to compute the driving current of the Integrate-and-Fire neuron. However, due to a different experimental conditions between the Human Embryonic Kidney (HEK) and the RGC experiments, there is a need for preliminary adjustment of the model before usage in the RGC context. Apart from the fact that both experiments have been conducted on different setups, the major difference between the two is the light source used for the photostimulation.

In the HEK experiment, a LED with a rather large spectrum was used. Its bandwith (Full Width at Half Maximum) is equal to $80\,\mathrm{nm}$. On the other hand, the light source used in the RGC setup has a much thiner bandwith of $15\,\mathrm{nm}$. Owing to the specific sensitivity of the ChrimsonR protein — see [128] for a comparison of the spectral sensitivities of most of the known channelrhodopsins, including Chrimson — two light sources with different spectra will have different impacts on the channelrhodopsin population, at constant light intensity.

To take into account this difference in experimental conditions, we ran preliminary analyses on voltage-clamp data performed on primate RGC. At this point, our first assumption is that all light-sensitive transition have the same spectral-sensitivity. This means that the values of the transition rates from all light-induced (orange) transitions in Fig. 13 (b) vary in the same way with respect to the wavelength of the stimulating photons. From a practical point of view, it means that we can apply an *Adjustment Factor* to the stimulus light intensity in order to take into account the shift in spectral content of the light source.

Our second hypothesis is that the intrinsic behavior of ChrimsonR is independent of the host organism. This claim is supported (on other channelrhodopsins) by [94, 105] and is the reason why we rely on electrophysiological experiments performed on HEK cells. Obviously, some factor change from one host to the next, starting with the total number of ChrimsonR proteins expressed within the host and which dictates the total amplitude of the responses, as well as the expression of molecular material specific to the host (e.g. voltage-gated ion channels) that can induce additional phenomena. This is why, (i) our analyses focused on the features of the response that are intrinsic to the single channel: its ON dynamics and its activation curve, and (ii) the voltage-clamp experiments were carried out in conditions where the *active* channels of the neuron are silenced: voltage clamped to $-60\,\mathrm{mV}$. Note that both features on which the analysis relies are aspects of the ChrimsonR behavior which are well captured by our Markov kinetic model.

The procedure is complicated by the existence of the slow adaptation mechanism that we uncovered in the previous chapter. Due to the reduced time laps between two successive stimulation blocks, the initial conditions are uncertain. This adds a source of ambiguity in estimating the Adjustment Factor. Moreover, the time interval between subsequent stimulation blocks was not recorded, which further complicates the procedure. However, we know when stimulation blocks have been presented subsequently, and that no light stimulation has been applied between the blocks. We therefore reconstruct a single stimulation pattern by setting an arbitrary resting period of a few seconds between the blocks.

In order to assess the value for the Adjustment Factor — mathematically defined as a multiplicative factor that we apply to the light stimulation values before performing the simulation of the response of the ChrimsonR model — we performed a least-square search on (i) the Adjustment Factor, (ii) the initial conditions (mostly the fraction of channels which are stuck in the slow state S at the beginning of the first block) and (iii) the expression level. We ran this procedure on one data set with two stimulation blocks with different values for the light intensity (1.6×10^{16} and $3.15 \times 10^{17}\,\mathrm{ph\,s^{-1}\,cm^{-2}}$). The optimum triplet for the factor, initial conditions and expression level yield the factor that we should apply, after verification that the initial conditions provided by the procedure lie in a realistic range (fraction of channels in slow state S between 0 and 95%).

We ran this procedure for three parameter sets from Chapter 3, which provided us with one factor associated with each set (the color code is kept consistent with Fig. 26 from Chapter 3: 🟥 Cell 2 / 🟦 Cell 3 / 🟩 Cell 5).

The Adjustment Factor was then validated by comparing the simulation of the model with voltage-clamp recordings from four additional cells. These recordings display a larger number of stimulation values and a wider range of stimulation patterns. Figures are shown in Appendix A.

Description of the raw data : main observations

The stimuli come in two main forms : (i) single square pulses of the same amplitude and of increasing duration (ranging from approx. 20 ms to approx. 4 s) and interspaced by time intervals ranging from 1 s to 2 s and (ii) sequences of ten square pulses of the same amplitude and with increasing frequency (ranging from approx. 2 Hz to approx. 25 Hz) and a duty cycle of 50%. The amplitude of the pulses could take two values : 1.6×10^{16} and

$3.15 \times 10^{17}\,\mathrm{ph\,s^{-1}\,cm^{-2}}$. The two forms of stimuli are shown in Fig. 28, along with sample recorded spike trains.

The action potentials were recorded using the cell-attached technique. In this method, the recording pipette is tightly sealed to the cell membrane but the membrane remains intact. Therefore, we do not have access to the membrane potential of the neuron, but the action potentials triggered by the recorded cell evoke a large impulse in the recordings and there precise timing can easily be recovered from the raw data using a simple thresholding method. In a few recordings, we had to correct for the drift in baseline before applying the thresholding method. In these cases, the baseline correction was performed by fitting the baseline using a linear regression and then subtracting the estimated drift from the signal.

From now on, we consider that each experimental block is the data of a stimulus (time and light intensity) and a sequence of spike times. In Fig. 28 to Fig. 30, we display the data obtained for a single RGC to the two types of stimuli in three different ways. When applied to the same data, these three methods reveal different aspects of the temporal organization of the spike trains.

In the first method, termed *spike count*, we plot the total number of action potentials which have occurred since the beginning of the experiment. This function is obviously increasing. It is very similar to a single line of a raster plot, i.e. we immediately see when the spikes are happening, but it also quickly reveals other phenomena such as the difference in spike rate between stimulation and spontaneous activity through comparison of the slopes of the spike count. This method is shown in Fig. 28a and Fig. 28b.

In the second method — Fig. 29a and Fig. 29b — termed *Interspike*, the spike train is represented as a scatter plot. For each spike in the spike train, a dot is displayed. Its abscissa is the absolute time of the spike and its ordinate is the time since the previous spike. This method is particularly useful to observe the local dynamics of the response during time intervals of increased response, i.e. during the stimulation pulses. Compared to the *spike count* representation, it reveals several important phenomena. First, it allows a quicker comparison between the firing rates evoked for the different pulses, by looking directly on the y-axis. A smaller ordinate means that the interspike interval is shorter, and thus that the local firing rate is higher. Consequently, this representation enables a direct assessment of slow adaptation phenomena. It also reveals the quicker adaptation phenomena, through the overall shape of the within-pulse interspike interval. Finally, it highlights the noise level through the variability in interspike duration.

The third method is a classical method called *Spike Density Function*[131]. It is obtained by convolving the spike train with a kernel function. Here we choose a centered Gaussian kernel, with a variance adjusted by hand. This method enables a rapid and efficient estimation of the firing rate and works even if we are only provided with a single trial. This method is pictured in Fig. 30a and Fig. 30b.

Simulating the model

Following the notations from [113], and as for any Integrate-and-Fire model, the spiking behavior of the neuron is performed by numerical simulation of the differential equation ruling the subthreshold membrane potential:

$$dV = \left[-\frac{1}{\tau}(V - V_{leak}) + I_{stim}(t) + I_{hist}(t) \right] dt + \sigma \, dW_t \quad (13)$$

where V is the membrane potential. I_{stim} is directly computed as the conductance of ChrimsonR (through simulation of the model introduced in the previous chapter) multiplied by a factor termed *Expression Level* representing the total number of proteins expressed in the membrane of the host cell. I_{hist} is a feedback and is a sum of stereotyped currents that are added for each spike triggered. In practice, we made the choice that the shape of this current be a simple decreasing exponential and we estimate the amplitude and time constant of this term through the global fitting procedure. V_{leak} is the leak potential and τ the time constant of the membrane. σ is the noise level and W_t is a Wiener process.

Finally, whenever the membrane potential V reaches the threshold $V_{th} = 1$, a spike is triggered and V is reset to V_{reset} which is set to 0 as in [113].

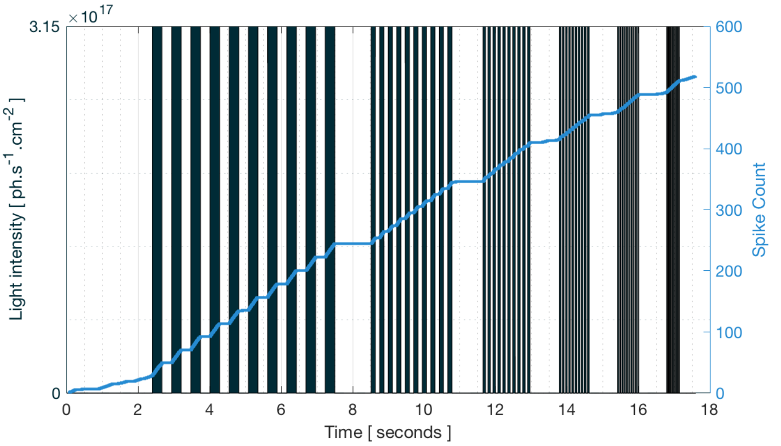

(a) Spike Count on Set #1

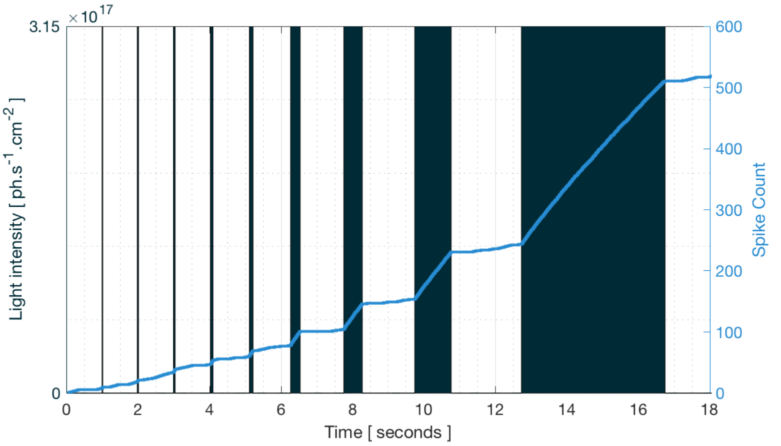

(b) Spike Count on Set #2

Figure 28: **Raw cell-attached data**: the data from two stimulation blocks displayed using the spike count method, where each spike increments the value of the curve at the time of the given spike. Top: a sample spike train in response to the first type of stimulation block (frequency patterns). Bottom: the response to the second type of stimulation pattern, square pulses with increasing duration.

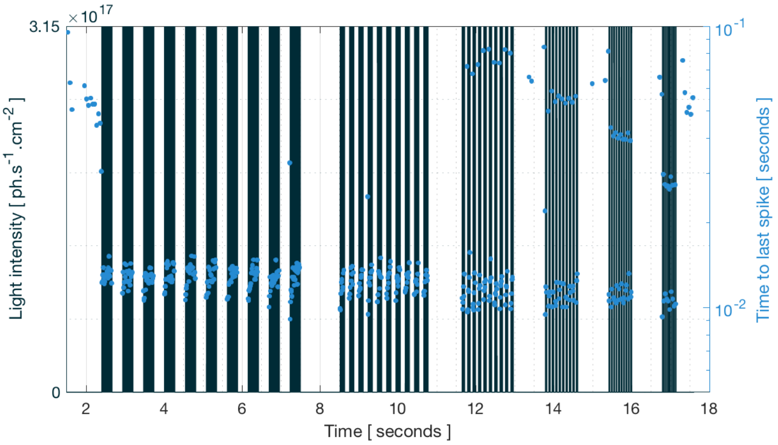

(a) Interspike on Set #1

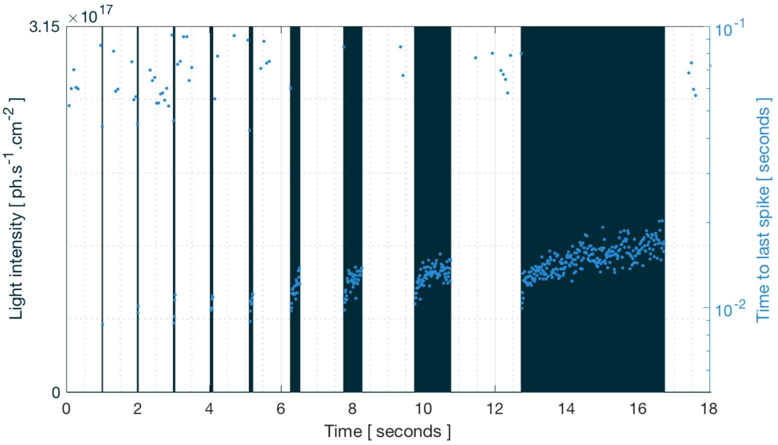

(b) Interspike on Set #2

Figure 29: **Raw cell-attached data**: the data from two stimulation blocks displayed using the Interspike method, where the coordinates for each dot are the spike time and the time interval since the last spike.

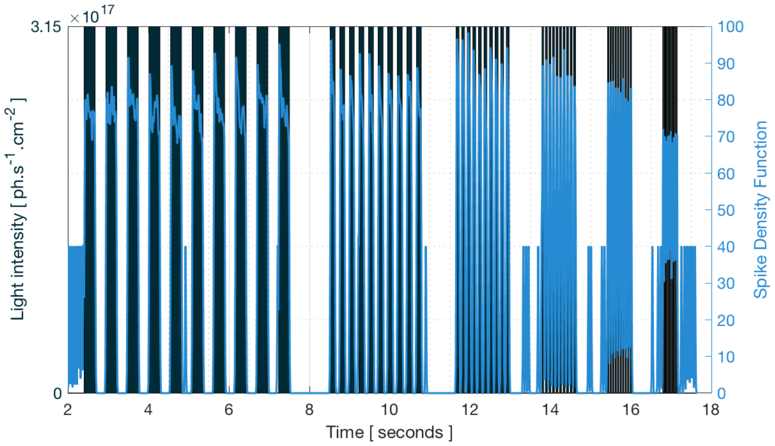

(a) Spike Density Function on Set #1

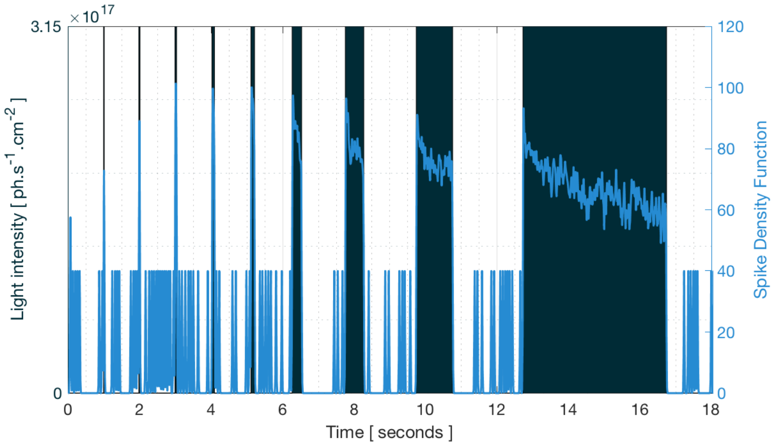

(b) Spike Density Function on Set #2

Figure 30: **Raw cell-attached data**: the data from two stimulation blocks displayed using the *Spike Density Function* method, where the spike train is convolved with a Gaussian kernel. This method yields an estimate of the instantaneous firing rate based on a single trial.

Computing the likelihood

The objective function we use to estimate the parameters of the L-LN-IF model is the likelihood that the model generated the observed spike train. Following the method described in [108], we carried out the numerical simulation of the Fokker-Planck equation describing the evolution of the probability density P of the subthreshold membrane potential:

$$\frac{\partial P(V,t)}{\partial t} = \frac{\sigma^2}{2}\frac{\partial^2 P}{\partial V^2} + \frac{1}{\tau}\frac{\partial[(V - V_{rest})P]}{\partial V} \tag{14}$$

where $V_{rest}(t)$ is defined as the stationary point of the noiseless subthreshold dynamics of Equation 13. The probability density is reset to a Dirac in the reset potential V_{reset} at the time of each *observed* action potential. Additionally, the feedback current was also computed based on the timings of the observed action potentials. The leak of probability density through the threshold potential V_{th} is the instantaneous firing probability density.

Estimation procedure

The first step of our estimation procedure consist in finding acceptable initial condition for ChrimsonR (as in the *Comparison with HEK293 voltage-clamp data* section, mostly the initial fraction of channels *blocked* in the slow state S) through a least-square minimization procedure similar to the one performed to estimate the Adjustment Factor. Intead of minimizing the difference between the recorded and simulated conductances (or currents), we use the Spike Density Function as the target. As shown in Fig. 32, the two curves have very similar shapes. The Adjustment Factor was kept fixed to its value established based on voltage-clamp data, and the expression level was left free in order to cope with the difference in unit between ChrimsonR conductance and the Spike-Density Function (SDF).

This procedure was run for each dataset and for each parameter set used for the protein. The initial conditions were then fixed to the value provided by the least-square procedure and used to estimate the parameters of the spike generation mechanism.

We then looked for a starting point for the likelihood maximization procedure. The starting point was fixed by choosing biologically realistic values, found in the literature, for parameters such as the time constant of the cell membrane and the temporal extent of the feedback contribution, and by adjusting manually the other parameters (expression level, amplitude of the feedback contribution, amplitude of the noise and resting potential) by having the simulated spike train match the recorded one. The matching criteria were (i) the total number of spikes, (ii) the precise relationship between spiking rate and stimulus intensity (observed on the *interspike* figure), (iii) the level of the spontaneous activity and (iv) the noise level.

Finally, a gradient ascent with optimal step was performed on the likelihood until convergence. Where convergence was defined as a threshold on the step taken, either in the parameter space or in the value space.

Software implementation

The software for simulation, likelihood estimation and likelihood maximization consists in a Matlab part and a C++ part. The Matlab part could implement the simulation part which is not too demanding computationally, but was mostly used to generate figures based on data output by the C++ part. The C++ part allowed to speed up the computation of the likelihood. It relied on the Armadillo library[125] for linear algebra computations, and took advantage of the fact that most of the computation of the likelihood can be cut into inter-spike chunks, and is therefore heavily parallelizable. Parallelization was implemented using the Intel *Threading Building Block* library.

The C++ program was compiled and run on a machine running on CentOS and equipped with 256 Intel XeonPhi cores.

4.3 RESULTS

Scaling factor in light intensity dependence

The analysis of the light intensity scaling factor yielded numerical values which are much higher than what we have been expected. We found a value of 12 for the most sensitive parameter set from Chapter 3 (the blue curve on Fig. 26) and 37 for the other parameter sets (cyan and magenta curves) whose activation curve lies with the ones from the three remaining parameter sets. The known differences in spectral content along with prior knowledge on the sensitivity spectrum of the protein[128] should lead to an Adjustment Factor in the order of 2.

Fig. 31 shows the simulation for all three parameter sets with the Adjustment Factors provided by the least-square procedure. Additional figures validating of the procedure are shown in Appendix A. The display of the responses to the full stimulation blocks show that the amplitude ratio between the two stimulation values is well respected and the four zooming windows show that the onset dynamics are also well reproduced.

Careful calibration of the light sources from the setup of both experiments have been performed anew, in order to discard the eventuality of a mistake in the calibration procedure. Both calibration procedures yielded results close to the original ones.

If not coming from the calibration procedure, the discrepancy between the two sources of observations could be caused by the change in light source, and more precisely by the change in spectrum. That would challenge our initial assumption that all light-induced transitions have similar wavelength sensitivities.

To date, we do not disqualify the possibility that this discrepancy is due to an unanticipated biological phenomenon that we fell to take into account in both the experiments and the modeling.

Fitting the initial condition using the Spike Density Function

Fig. 32 shows the result of the least-square fitting procedure for the initial conditions. In fact, since these two datasets come from subsequent block from the same cell, the initial conditions where only fitting on the first block (Fig. 32a). The second block was used as a validation (Fig. 32b).

Validation of the software

In order to validate the software, we performed the optimization procedure individually on the two stimulation blocks. Monitoring of the optimization procedure itself is shown in Fig. 33. We display the evolution of the log-likelihood with the number of iteration in the search algorithm, as well as the gain in log-likelihood that is made at each step. Both figures show that a plateau is quickly reached after a few iterations.

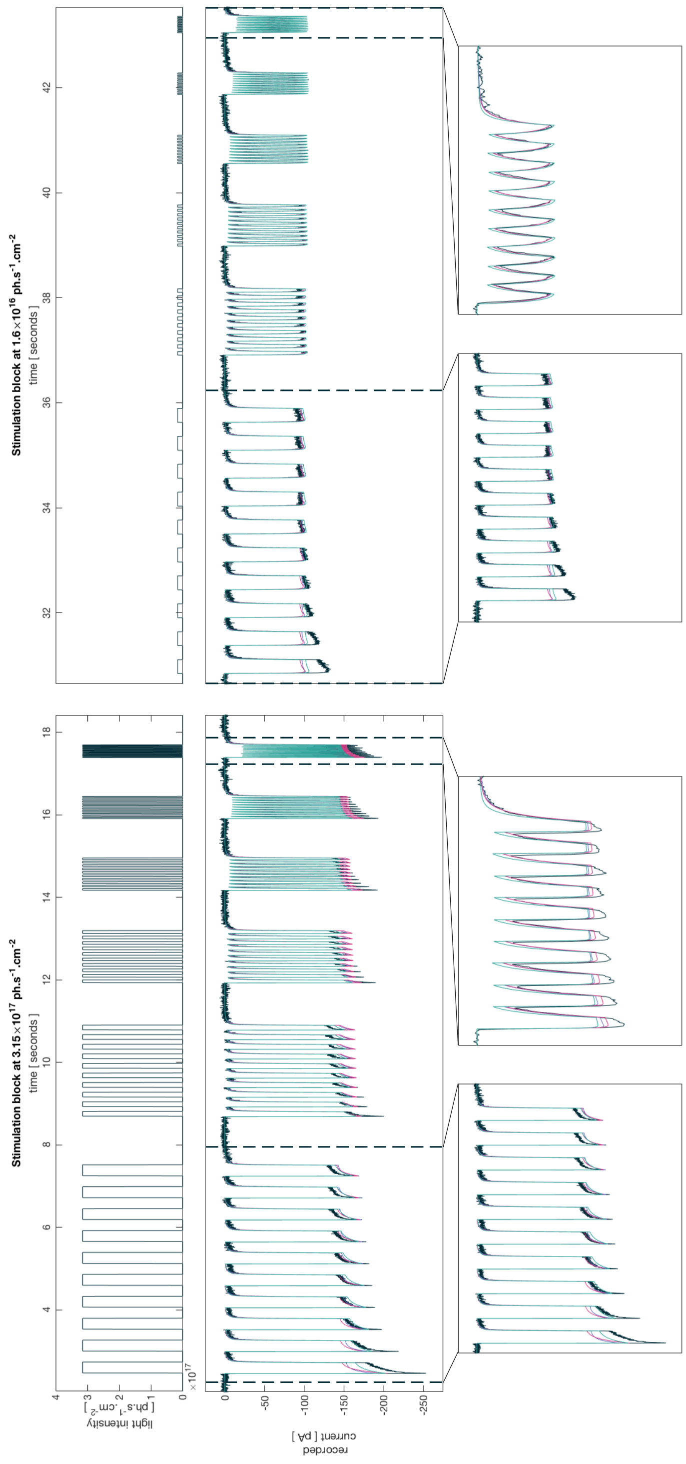

Figure 31: **Result of the least-square minimization procedure** leading to the estimation of the Adjustment Factors for the three different parameter sets. Even if presented in separate parts of the figure, the simulation or the response of the protein model was run in one piece for both stimulation blocks. The two blocks (which are known to have been presented successively) were arbitrary separated by a 5 s interval.

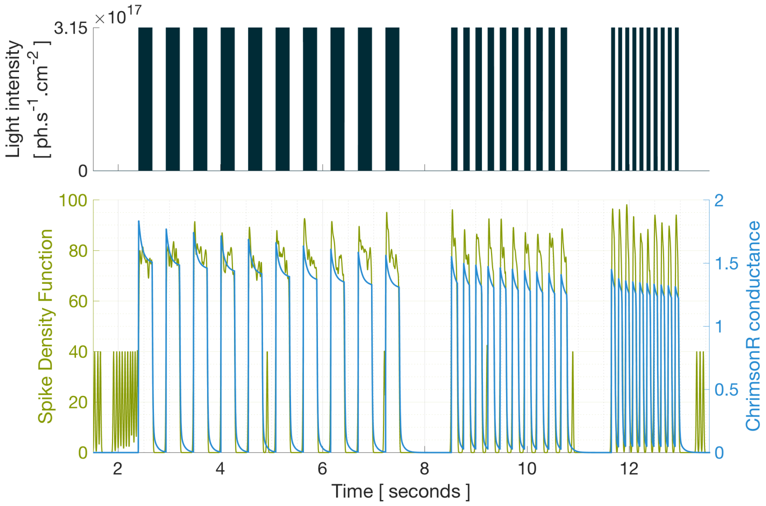

(a)

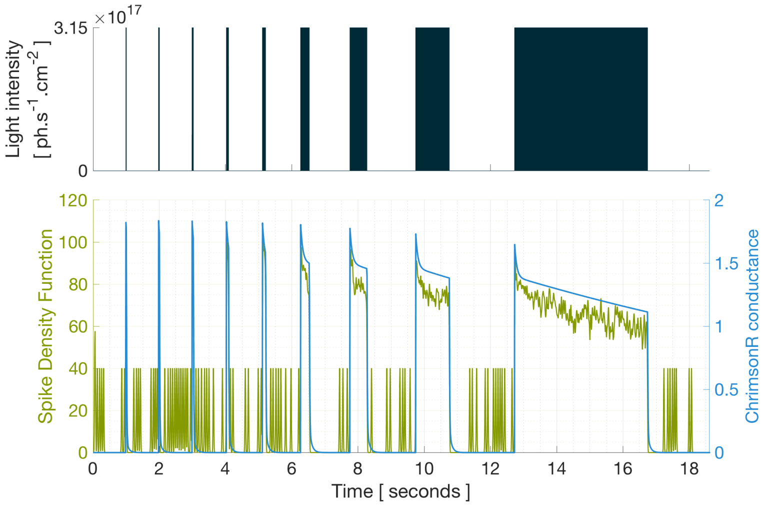

(b)

Figure 32: Fitting ChrimsonR's initial conditions using the Spike Density Function: results on two data sets (for the same parameter set for ChrimsonR)

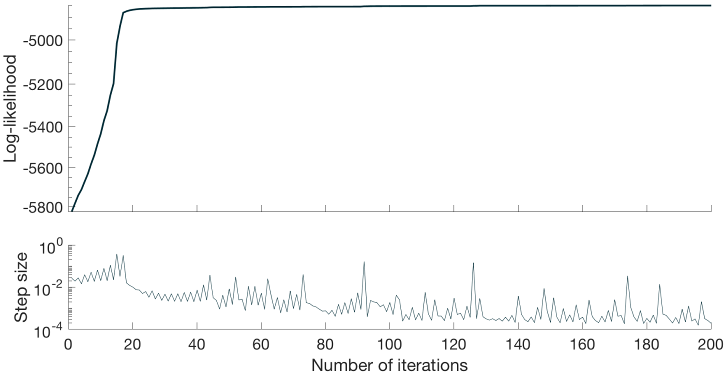

(a) Optimization monitoring on Set #1

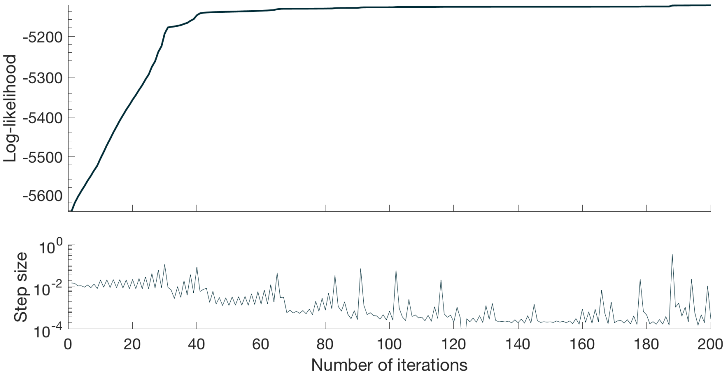

(b) Optimization monitoring on Set #2

Figure 33: **Monitoring the progression of the optimization procedure**: The top curve shows the evolution of the log-likelihood with the number of iterations in the procedure. The two subfigures show two independent optimization runs launched on two different sets recorded with the same cell.

PARAMETERS	SET #1	SET #2
k_{ChR}	4.95×10^{-5}	3.95×10^{-5}
σ	1.2×10^{-3}	7.07×10^{-4}
τ [ms]	9.4	16.7
V_{leak}	0.83	0.91

Table 4: **Numerical values resulting from the optimization procedure**: the results are shown for both datasets (same cell, different stimulation patterns).

ASSOCIATED TIME CST	SET #1	SET #2
$\tau_1 = 200\,ms$	-0.35	-0.33
$\tau_2 = 100\,ms$	0.21	0.28
$\tau_3 = 50\,ms$	-0.17	-4×10^{-3}

Table 5: **Feedback kernel amplitudes**. Note that in this version of the code, instead of estimating the amplitude and time constant of a single exponential term kernel, we fitted the amplitudes of a linear combination of three terms with fixed time constants around the expected value.

With the parameter sets resulting from these procedure, we performed a simulation of the spiking behaviour for each set. The results of these simulations are shown in Fig. 34 and Fig. 35 (for sets 1 and 2 respectively). Overall, there is a very nice adequacy between the recorded spike train (shown in green) and the simulated spike train (in red). From the *Interspike* comparison of the two spike trains, it seems that the noise level may be slightly overestimated by the fitting procedure. This is probably due to the fact that the model is not well adapted to capture both the spontaneous activity (noise in absence of stimulation) and the noise level during stimulation. Such characteristics, if they were to remain throughout the whole set of fitted recordings, would lead to an underestimation of the encoding capacity of the BCI.

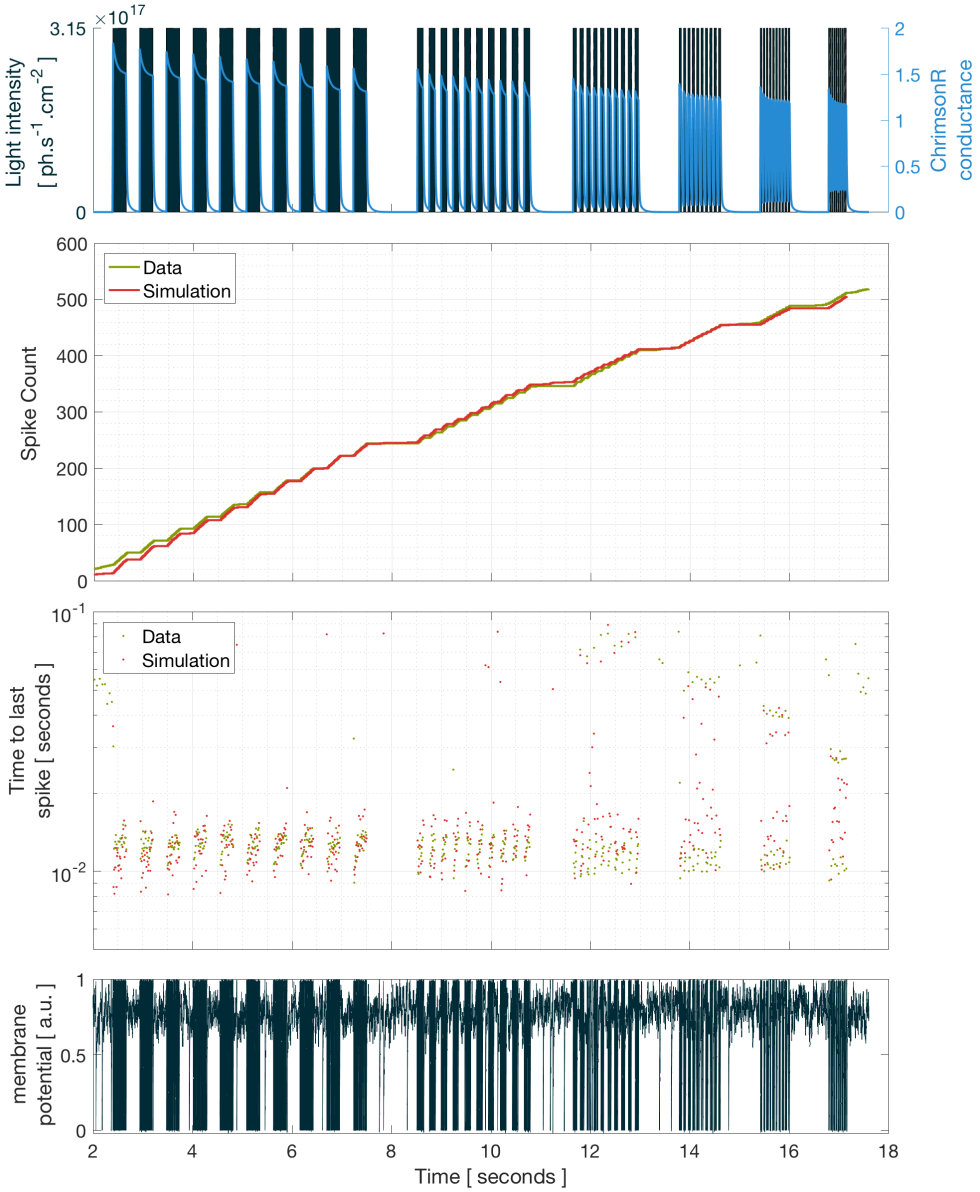

Figure 34: **Simulation of the spiking response of the model with the parameter set resulting from the maximization of the likelihood**. Top row: the stimulus (left axis) and the simulation of the ChrimsonR model (right axis). Second and third rows: comparison of the simulated spike train (red) with the recorded spike train (green). The spike trains are displayed with the *spike count* and *interspike* methods. Bottom row: the evolution of the simulated subthreshold membrane potential leading to the spike train.

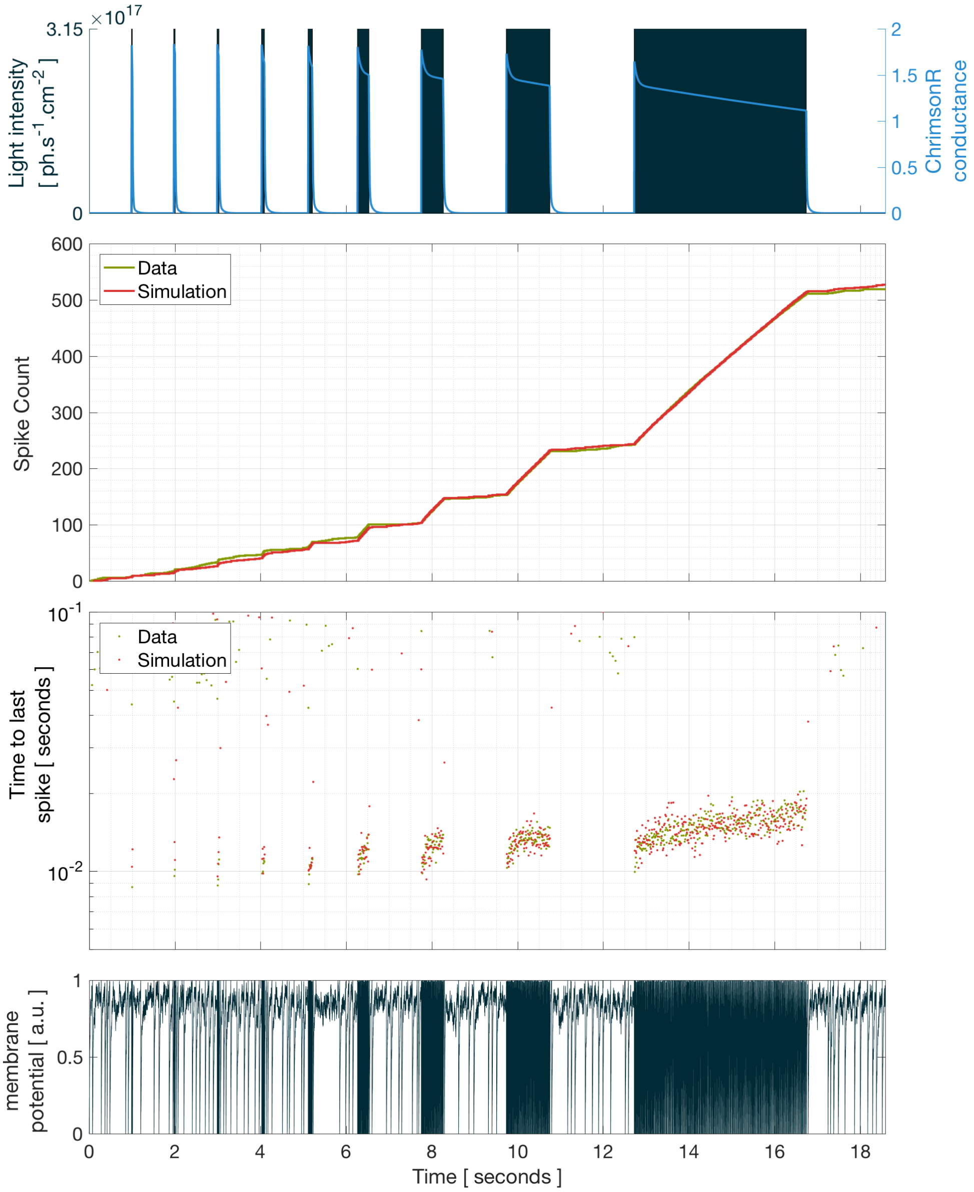

Figure 35: Same as in Fig. 34, but showing the result of the optimization procedure on another dataset from the same cell.

4.4 FUTURE WORK

In this chapter, we provided the proof-of-concept for the development of a model describing the spiking generation process of the primate ChrimsonR-expressing Midget Retinal Ganglion Cells. Preliminary results suggest that the model is able to capture all key features of the spike trains temporal organization — including the precise timing and jitter of single spikes — and the spontaneous activity.

This leads the way to the complete fitting of a great number of cell-attached recordings of primate Midget RGC. To be completely satisfactory, the fitting procedure will have to be run on longer stimulation protocols, including possibly three unit blocks such as the ones shown in Fig. 34 and Fig. 35, and involve several light intensities. Regarding the fact of using several light intensities in the same stimulation protocol, this will create an additional challenge during the step where the initial conditions for the ChrimsonR model are assessed. In fact, due to the nonlinear relationship between the conductance of ChrimsonR and the instantaneous firing rate, we will have to guess the shape of this relationship in order to specify properly the cost function defining the distance between the conductance and the Spike Density Function.

However, given that the shapes of the conductance curve and Spike Density Function and the fact that we were able to obtain very satisfactory fits with the preliminary datasets, the chances that the framework quickly provides good results on extended datasets are high.

4.5 DISCUSSION & CONCLUSION

Since the model relies on the model from Chapter 3, it inherits its strengths and weaknesses. one of the most remarkable feature of the model is that it scales well with light intensity. This is particularly relevant if we want to use this model as a tool in the conception of stimulation strategies for the GS030 project. This is usually a weakness of this kind of model, since it usually primarily relies on a linear filter to describe the relationship with the stimulus[113]. However, we manage to avoid this pitfall because our model for ChrimsonR captures well its response at all intensity values.

Additionally, the model relies on few parameters, which reduces the risk of *overfitting* (the fact that the model and associated parameter set are particularly well adapted to the data, but have a weak predicting power for other situations or data). Moreover, the parameters of the model have a clear meaning, allowing for an easy interpretation of the model and of possible differences from one parameter set to the other. It also allows more flexible handling of the model in a speculative setting, such as experimenting the consequences of a different level in noise.

The model obviously inherits the limitations of the ChrimsonR model as well, and primarily its lower predicting power at very high dynamics (frequency >100 Hz).

Another limitation is the departure of the model from the true biological mechanisms underlying spike generation. As highlighted by Jolivet et al., this will constrain the use of the model to situations similar to the ones encountered during the fitting procedure. This aspect should have a limited impact since the stimulation protocols were specifically designed to probe situations that would occur with the actual stimulation patterns.

For instance, in some cells subject to the same stimulation patterns, *excitation block* has been observed at $3.15 \times 10^{17}\,\mathrm{ph\,s^{-1}\,cm^{-2}}$. We deliberately ignored these cells in our modeling effort. We made the choice to develop a model capturing the spiking behavior of RGC in regular conditions. However, *excitation block* is a situation that we absolutely want to avoid — at least for a large fraction of the transfected cell population — since it is a situation where the high intensity stimulus results in the absence of spikes resulting in downstream structures receiving a message meaning the opposite of the intended one. Therefore, there is a need for either a more complex model which is able to capture the excitation block events as well as the correct firing patterns outside these events or a complementary model able to predict the occurrence of excitation block and thus informing when the simpler model is valid.

However, once properly fitted using a significant number of datasets, the model will already offer a great number of possibilities to further study the BCI.

The first possibility is the study of the temporal resolution. This could be achieved by having a bank of high temporal resolution stimulation signals and observe the repeated responses of the model to the full-resolution stimulus as well as to degraded versions of the stimulus (with lower temporal resolution). Analysis of the raster plots and Peri-Stimulus Time Histograms will reveal wether the spike trains encode the fine temporal details of the original stimulus or if some is lost in the encoding process.

The model can also be used *as is* to simulate the spiking responses to natural stimulation patterns with the strategies that are currently under investigation for the first clinical trials. Examination of the raster plots and Peri-Stimulus Time Histograms may give a first valuable feedback regarding the validity of the candidate strategies. There is already a vast literature on the topic in the *Retinal Implant* community where the best stimulation patterns are studied in order to evoke the most precise perceptions[69, 70, 71]. The model presented here the tool which is perfectly adapted to start such effort in the *Optogenetics* field. Algorithms taking advantage of a modeling step are already available[6] and can guide us through this work.

In this line of work, a fruitful amelioration would consist in working on a decoding algorithm that would be able to perform the inverse work and infer the stimulus from a spike train. In their original paper, Pillow et al. already show how it is possible to discriminate between two stimuli with the L-LN-IF model. Full decoding is a more complex task. A fraction of the authors published an additional study addressing specifically this decoding issue[114].

The combination of the encoding model and the decoding algorithm is a very powerful tool for the development of stimulation strategies. One of the main benefits is that we can show an healthy subject a video that get close to what we might expect the targeted RP patients to perceive. Even if an RP patient will probably perceive the stimulation in a different way, owing to the unnatural stimulation pathway, the healthy subject would perceive a film which is close in terms of transmitted information. Careful evaluation of what the healthy subject can and cannot do will be a good clue of what we may the expect the RP patient do be able to do.

A similar setting can be used to discriminate between several types of visual information that can be displayed to the patient. For instance, we could consider projecting contour information or depth information on the patient's retina. Using the model-decoding algorithm combination, we could ask healthy subjects to perform a variety of perceptive and motor tasks and thus discriminate the most relevant information based on the results of the subject.

The work presented in this chapter offers the great promise of having a model describing precisely the BCI of the GS030 project. This is an invaluable tool to study the rate of information that can be transmitted to the RP patients through the light projectors and the transfected retina. This in turn can be used to draw preliminary conclusions on the level of vision restoration that we can expect from the device.

Additionally, the model is central in addressing a number of other issues related to the interface such as the type of information which is actually displayed, the encoding strategy which needs to monitor the tradeoff between the maximization of the current information transfer rate and the actionability of the channelrhodopsin population (with a minimal fraction of proteins being stuck in the slow state S).

Part IV

DISCUSSION AND CONCLUSIONS

5

DISCUSSION

The purpose of the stimulation strategy is to provide a full processing pipeline, from the sensor to the interface, such that — after an appropriate rehabilitation phase — the elicited neural activity is a faithful representation of the environment and can be interpreted by the CNS in order to guide behavior in an intuitive and effortless fashion.

Main challenges

Natural scenes contain large amounts of information, therefore the success of a stimulation strategy primarily relies on its ability to transmit large amounts of information through the BCI. This is the main blocking issue in retinal prostheses implementing the interface with a retinal implant and in Sensory Substitution strategies. The GS030 project addresses this issue by using the Optogenetics technology, which offers cellular spatial resolution as well as high temporal resolution. The whole processing pipeline, including the sensor and the light projector, supports a millisecond resolution of the visual information processing[55].

In addition to optimizing the capacity of the interface, a preprocessing step simplifying the content of the visual information displayed to the patient will help make the transition from the lab to the real world. The function of this step is to compensate for the limitations of the Brain-Computer Interface by only displaying the information most relevant to the patient.

Additionally, we argue that the success of the whole project relies both on problems that depend specifically on the condition of the RP patient and on problems which are related to the device and to the stimulation strategy but not to the disease itself. In Section 5.2, we list a number of issues which are related to the device and not to the diseases. Because of the complexity of the task, we argue that these problems should be carefully addressed and quantified before the system is tested on patients suffering from RP.

The issues specifically related to the disease are issues that cannot be addressed before testing the system with RP patients. They remain the major source of uncertainty regarding the success of the project. These questions are discussed in the conclusion of the document.

Structure of the stimulation strategy

The information processing pipeline consists in (i) a sensor, (ii) a processing step and (iii) a Brain-Computer Interface. The sensor imposes the way raw data is represented, and the BCI imposes the transformation between a signal that we can control (the light level for each pixel of the projector in our situation) and the nervous activity that this signal induces.

The information that we can communicate to the retina through the BCI is necessarily a unidimensional time-varying signal for each pixel of the projector. It would be illusory to try to stimulate the different RGC types with different types of information. In fact, this would require having a good knowledge of the types of all ganglion cells in the stimulated region and being able to continuously track their location as the eyes move. Therefore, the processing step necessarily consists in two steps: a first step extracting the relevant information from the output of the sensor and a second step that controls the projector so that the extracted information is communicated to the nervous system in an efficient manner. These two steps are called *pre-processing* and *mapping* in Fig. 36.

The structure for the stimulation strategy that we propose here goes against a more natural strategy which would consist in predicting the spike trains that would occur in natural vision and try to elicit spike trains as close as possible from the prediction. The two main reasons why reproducing *natural vision* is not a viable option are (i) the complexity of the signal that we should be reproducing and (ii) our current limited understanding of natural vision. In fact, the stimulation of the retina occurs at the level of the Retinal Ganglion Cell layer. Classification of the RGC in mammals is not complete yet, but it is estimated that there are at least 30 different types[126]. Each RGC type encodes a specific feature of the visual scene and the population of a given type is arranged according to a regularly spaced mosaic through the entire retina. Mosaics from different types are interwined. Re-

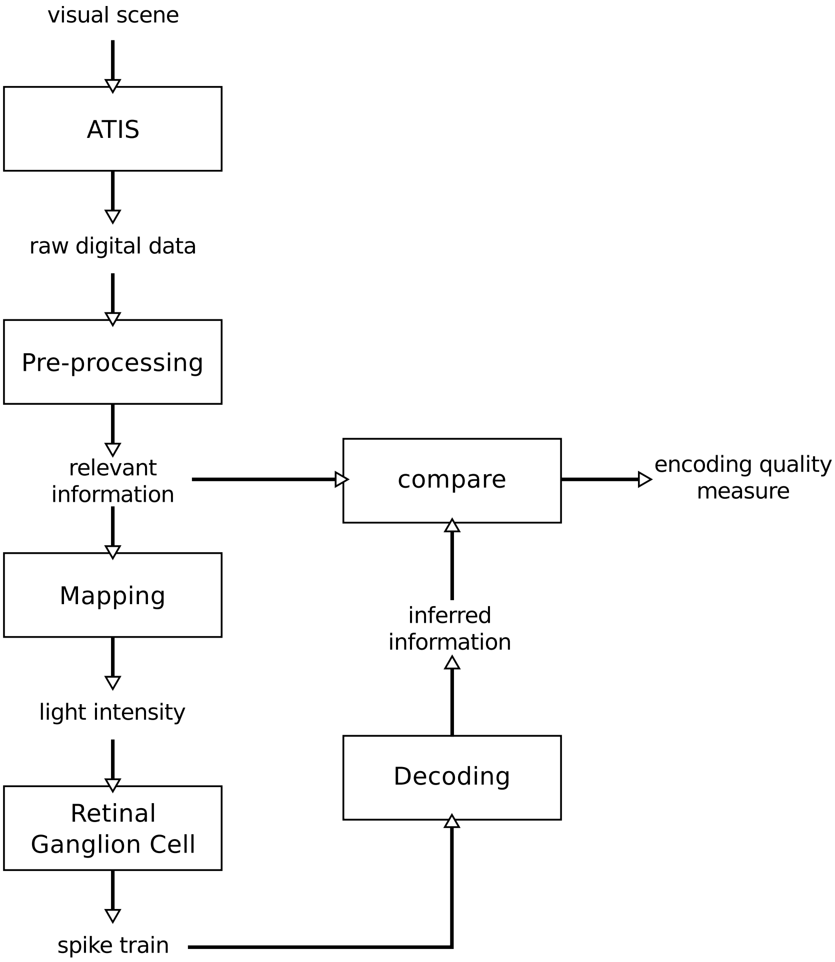

Figure 36: **Structure of the stimulation strategy** The visual information undergoes a number of transformations between the real scene and the patient's nervous system. First, the photons carrying the information are captured by the sensor and stored under a digital form. This representation of the raw data is not necessarily adapted for immediate communication through the interface. A pre-processing step is therefore required to extract relevant information and set it in a representation adapted for further communication. This information will take the form of a two-dimensional array of time-varying signals encoding a specific feature of the visual scene. The function of the *mapping* step is then to control the projector so that the spike trains elicited in the Ganglion cells faithfully encode the original features.

producing the natural behavior of the ganglion cell layer thus requires several complicated steps. First, it requires to be able to predict the expected spike train for each type. Even if some cell types are very well studied, we do not possess a satisfactory model for all cell types. Moreover, that would then mean that we have to identify the type for each cell in the retina of the patient and that we are able to track it and adapt the stimulation according to the movements of the patient's eye.

Even if the strategy is adapted to a single RGC type, we argue that this strategy is formally equivalent to the one that we introduced. The feature is implicitly chosen based on the feature naturally encoded by the selected RGC type, and the mapping step is carried out because ganglion cells naturally adapt their response to encode efficiently the visual features in their spike trains. The problem therefore relies in all the implicit choices that are made, and the subsequent difficulty to assess the benefits or issues associated with each step of the processing pipeline.

The major benefit is that it allows addressing the two problems, (i) selecting the content of the transmitted information and (ii) maximizing the capacity of the interface, separately and solving them independently. In particular, the question of the information that should be communicated to the patient given the limited rate of information that can flow through the interface can be addressed based on psychophysical studies conducted with healthy subject and a Virtual Reality interface mimicking the degraded interface. We discuss the content of such studies in Section 5.2.

Dealing with spikes

One of the main challenges of the project — as far as information processing is concerned — is the interpretation of spike trains, which are the support for information processing and communication in the CNS. The difficulty is twofold. First, we are not used to carry out computations in a spike-based manner and therefore our understanding of the way information is encoded into spike trains is limited[51]. Second, we lack tools to predict how the brain will interpret the artificial spike trains triggered by the retinal prosthesis.

Once again, we encourage separating these two questions and address them independently. The first issue can be addressed through mathematical analysis, and we discuss this work in Section 5.1. The second question exceeds our current understanding of the brain. We discuss this question in the conclusion, it belongs to the category of questions that can only be addressed when testing the full system on RP patients in a real setting.

5.1 OPTIMIZING THE BRAIN-COMPUTER INTERFACE

The question of the rate of information that can be transmitted through the interface is central to our project. Characterizing its properties and making sure that it is used to its full potential are key to the success of the entire project. This work is carried out by studying the relationship between the signal representing the information that we want to display to the patient (the output of the preprocessing step) and the spike trains triggered as a result of the photostimulation.

The model that we developed in two steps in Chapter 3 and Chapter 4, which is able to predict the output spike train resulting from the stimulation by an arbitrary command signal, is the main tool in this line of work.

However, we are confronted with the difficulty of having to deal with information encoded in the form of spike trains. To overcome this issue, we can use two main approaches.

A simple solution consists in assuming that information is encoding into the spike trains according to a rate code, i.e. that most, if not all, information about the stimulus is contained in the firing rate of the neuron. Using the *Spike Density Function*[131] or a Peri-Stimulus Time Histogram (PSTH) on repeated responses to a single stimulus, we can transform spike trains into time-varying signals that we can analyze more easily. Because of its simplicity, this solution should be favored to start with.

A more ambitious solution consists in using the spiking model of Chapter 4 to build an inverse model reconstructing the original signal based on the spike train. As the previous one, this method is not concerned with how the brain actually interprets the information carried in the spike train. This method estimates more faithfully the content from the original signal that is actually encoded into the spike train. The major benefit of this method is that it provides a signal that can be directly compared with the original signal, without further correction or transformation.

Two main axes of research

The theoretical work on the BCI can be separated into two main axes.

The first one consists in studying the interface and characterizing its properties. The properties we are particularly interested in are its temporal resolution and its noise level. Estimating the temporal resolution of the encoding mechanism is essential because it directly affects the rate of information transmitted to the patient's nervous system. It also allows to adapt the temporal resolution of the signal output by the preprocessing step. In fact, trying to transmit more information (in particular with a temporal resolution that is too high) can only decrease the capacity of the interface.

The noise level is an important feature to study as well. If it is too high, then the signal at the level of a single pixel cannot be encoded faithfully into a spike train. In that case, compensatory strategies introducing redundancy between pixels should be considered. As in the case of the temporal resolution, the objective is to acquire a good understanding of the interface so that we know precisely what are its capacities and therefore have a precise idea of information that is actually accessible to the patient.

The second axis of research naturally follows the first one and consist in designing the mapping algorithm, taking into account all knowledge about the BCI, so that the interface can be used to its full capacity.

The questions that need to be addressed in priority are the level of the light source and the management of the slow adaptation process of ChrimsonR that we unveiled in Chapter 3. As we know, the projector mounted on the goggles worn by the patient consists in a light source and an array of micro-mirrors having two positions, ON and OFF[55]. The ON dynamics of the channelrhodopsin are dictated by the intensity of the light stimulus. Therefore, a higher light intensity leads to a higher temporal resolution for the interface, and following a higher capacity to transmit information.

However, higher stimulation of the channelrhodopsin population leads to more channels being stuck in the side reaction modeled by the slow state S. In turn, it leads to smaller responses from the channelrhodopsin population and therefore a lower excitability of the transfected cells, and a reduced capacity of the interface.

Moreover, before addressing these questions, we need to clarify the actual response range of ChrimsonR in their actual targeted environment, i.e. the Retinal Ganglion Cells of a living RP patient. We saw in Chapter 4 that moving from HEK cells to RGC and changing the light source imposed the use of an *adjustment factor* on the light intensity which is significantly bigger than what was expected based on the spectral sensitivity of the protein.

Monitoring the fraction of channelrhodopsins stuck in the slow state S is going to be an important aspect of the mapping step. In fact, due to the slow adaptation mechanism, excitable channels (i.e. the channels in states C_1 and C_2) should be considered as resources that can either be spent to trigger spikes or saved for later in order to preserve the level excitability of the population. The mapping algorithm must implement a strategy that is continuously solving this decision problem. Such algorithm can take two main forms. First, it can be implemented through a set of rules that ensure a minimum level of excitability at each time, for instance by imposing that the micro-mirrors are turned off at least 50% of the time on average. Such algorithm would be tested *a priori* with realistic stimulation patterns and making sure that the cell population reliably triggers spikes whenever it needs to. The second category of mapping strategies is slightly more complicated. It consists in using the five-state Markov model in order to continuously monitor the fraction of channels stuck in state S, and using this information to dynamically adapt the decision process to the current level of excitability of the population. This second type of approaches would allow more flexibility in the mapping process, and therefore result in an interface with extended capacity.

5.2 PREPROCESSING STEP

Why we need a preprocessing step

Existing vision restoration or substitution devices tend to show their limit when moving from the lab to the real world. Despite impressive results to a number of visually guided tasks, these systems haven't proved as useful as expected in improving the visual abilities of the patients in their everyday life. This phenomenon results from the great complexity of the visual scenes that can be found in the real world.

Simplifying the visual scene and selecting only the most relevant information can hopefully compensate for this problem, and still provide useful information to the patient.

The preprocessing algorithm primarily needs to adapt the temporal and spatial scale of the projected signal to the characteristics of the BCI. Additionally, the content of the signal can be chosen so that the signal reflects information that is directly useful to the the patient. Examples of features are brightness, contours, saliency or depth.

Using these features obviously requires real-time extraction from the sensor's output data, and that the algorithm can be implemented on the hardware embedded in the processor unit. This is an issue that we do not address here, but we discuss below some guidelines regarding the way they should be assessed and compared, provided that they are available.

Guidelines for testing Feature Extraction algorithms

The aim is to select the most valuable information so that it is useful to the patient despite the limited capacity of the BCI. The candidate algorithms should be tested using psychophysical studies with healthy patients with a Virtual Reality interface mimicking the BCI. The first requirement is therefore to reproduce a limited, noisy interface. Since the objective is to adapt the preprocessing algorithms to the actual BCI, the benefits of the study will increase as the simulation becomes more realistic. As in the characterization of the interface presented in Section 5.1, the development of an inverse model would be extremely useful in this context. It would allow to display to the healthy subject a movie faithfully reproducing the information contained in the predicted spike trains.

However, before we have a precise model (along with the inverse transformation) there are a lot of transformations that can affect the interface and model a limited interface. Possible means to limit the rate of information through the interface are (i) the spatial resolution, (ii) the temporal resolution, (iii) image distortion or (iv) the addition of different types of noise.

Since the aim is to provide a system that *compensates* for the complexity of the visual world, the tests should take place in a visually rich environment, and not in a lab setting where the scene is already simplified. However, because the aim is to be able to compare feature extraction algorithms, the tests should result in quantitative notation. Obviously, this is a difficulty because of the richness of the environment. However, settings such as the HomeLab or the StreetLab are perfect settings for this purpose where the patient / healthy subject can evolve in a complex environment which remains controlled and safe[124].

Using the Fourier algorithm in the framework

The algorithm presented in Chapter 2 extracts the spatial frequency content of the time-varying image. At the end of this step, the lower frequencies can be selected while the higher frequencies representing the finer spatial details of the scene are cut off. Application of the same algorithm (with minor adjustments) then converts back the frequency data into an image with simplified spatial details.

It is therefore a candidate algorithm for the preprocessing step. However, before being adopted, It should be tested according to the guidelines introduced in the previous paragraph and compared with other algorithms.

Important aspects of the device

Using the device will have a number of effects on visual perception which are inherent to the device, and not related to the condition of the RP patients it is designed for. These effects should be carefully studied before testing the system with RP patients so that we are sure that the problems that we address with these patients are specifically related to their condition.

Among the problems that can be studied with healthy subjects before the clinical trials are (i) the reduced spatial resolution of the camera and projector, (ii) the reduced field of view, (iii) the absence of color information and (iv) the delay induced by the processing steps between the sensor and the projector. All these problems come in addition to the limited capacity of the BCI. Their impact on vision should be minimized, or at least known.

6

CONCLUSION

Beyond the steps remaining before setting up a complete processing pipeline, there will still be a number of questions that can only be answered when the strategy is actually tested on patients suffering from RP. In this chapter, we make a non-exhaustive list of these questions and try to provide some intuition regarding the possible outcome of the actual observations.

Physiological state of the retina after photoreceptor degeneration

Using optogenetic therapy requires that some portion of the retina functional circuits remains intact. Work on the rd1 and rd10 mouse models of RP as well as in postmortem human retinas from RP patients suggest that this is often the case[96, 73]. However, the time course of retinal degeneration depends on the specific mutation causing the disease and there is important variability across subjects suffering from the same mutation. Therefore, the possibility to set up an optogenetic therapy for RP is real, but "the state of the retina of a particular patient at a given time cannot be predicted"[124]. Ultimately, individual examination of each patient based on *in vivo* imaging will be required to determine if the optogenetic strategy is appropriate.

Cerebral plasticity and homeostasis

Plasticity, the ability for the brain to modify its synaptic connections and therefore modify the computations carried out, is a very complex mechanism that is not uniform throughout the brain and over time. It is central to the computational power of the brain and to its ability to learn complex behaviors and adapt to new environment. However, plasticity is not everything and the co-occurring mechanism of *homeostasis*, the maintenance of a stable state, also plays a major role in the brain's function[50, 111, 33].

The observations in Sensory Substitution experiments that stimulation through a replacement modality is able to trigger responses in the areas originally responsible for the processing of information sent from the lost sense are very encouraging. It means that plasticity can take place on large scales and induce important changes in the way the CNS processes sensory information. However, and whatever the reasons, these experiments did not lead to the point where the subject can make satisfactory use of its substitution system. Consequently, it is a complex question to assess if its a problem with the interface or if there is an inherent limitation in the plasticity mechanism as well.

In cochlear implants, the age at which the procedure is performed is one of the primary factor of success. Beyond the age of 3.5, the chances of success tend to diminish and they significantly decrease after the age of 7[64]. The situation is different in our case because our targets are people who used to see. Consequently, they experienced the Critical Period and learned to see. This is good news because it means that the different structures devoted to visual perception are possibly functional and only need visual information to be able to perform their function. However, this does not mean that they are going to be able to adapt to the new type of incoming information.

Learning the new role of the Retinal Ganglion Cells

In natural vision, the different ganglion cell types send parallel streams of preprocessed information to downstream areas of the visual system. With the optogenetics-based retinal prosthesis, the spike trains of all ganglion cell types will encode information which has been set *a priori*. Even if we choose to favor one ganglion cell type and try to have this type reproduce the natural behavior, cells from all other types will respond to the stimulation (each one in its own way depending on its level of expression of the transgene and its own morphology) and the associated spike trains are going to travel to the brain.

Downstream visual areas receiving projections from both cell types must learn that the different populations are no longer encoding different features, as they used to in natural vision. Whether the downstream visual areas will be able to perform the operations that they used to carry out before the occurrence of the disease based on the new role of the ganglion cells is a very exciting question that only the clinical trials are going to answer.

Similarly, some areas receiving projections from the retina will receive a signal that is not relevant to the task they are supposed to perform. Will they be able to understand that the new signal is irrelevant, and that they should completely shut it off ?

Final comments

The GS030 approach to building retinal prostheses introduces new key elements, such as the millisecond temporal resolution of the entire processing pipeline and the optogenetics technology enabling single-cell spatial resolution. These novelties have the potential to break the barriers that have prevented previous attempts to restore vision to succeed, despite a few remaining questions which are left for the clinical trials with RP patients.

The work presented in this thesis participates in the effort to set up the processing pipeline from the ATIS camera to the light projector. It opens the way to a complete description of the Brain-Computer Interface, which will lead to a better understanding its information transmission capacity and allow the development of preprocessing algorithms adapting the content of the information displayed to the patients.

We are confident that this work will help address all the issues specifically related to the device and allow patients to fully benefit from the rehabilitation program, and hopefully recover vision in a way that makes a difference in their everyday life.

APPENDIX

A

SUPPLEMENTARY MATERIAL ON MODELING THE SPIKING BEHAVIOR OF TRANSFECTED CELLS

This appendix contains voltage-clamp data for four additional Retinal Ganglion Cells. They were used to validate that the factor used to make the transition between the HEK and RGC experiments were correct. Fig. 37 to Fig. 40 show the original data, superimposed with the simulations from the three parameter sets for ChrimsonR and optimized initial conditions. The Factor on the light intensity remained unchanged.

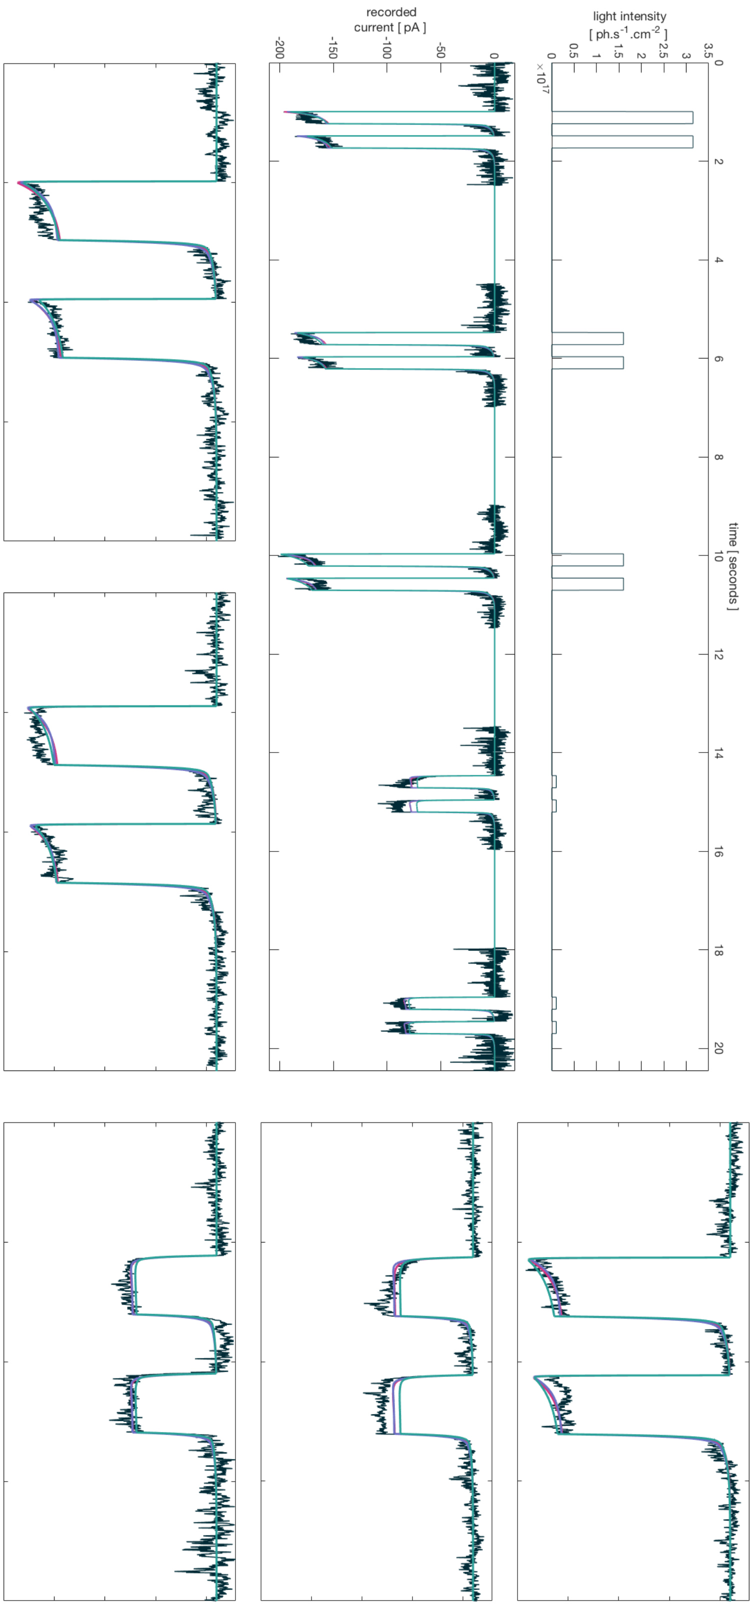

Figure 37: Validation on Cell #2

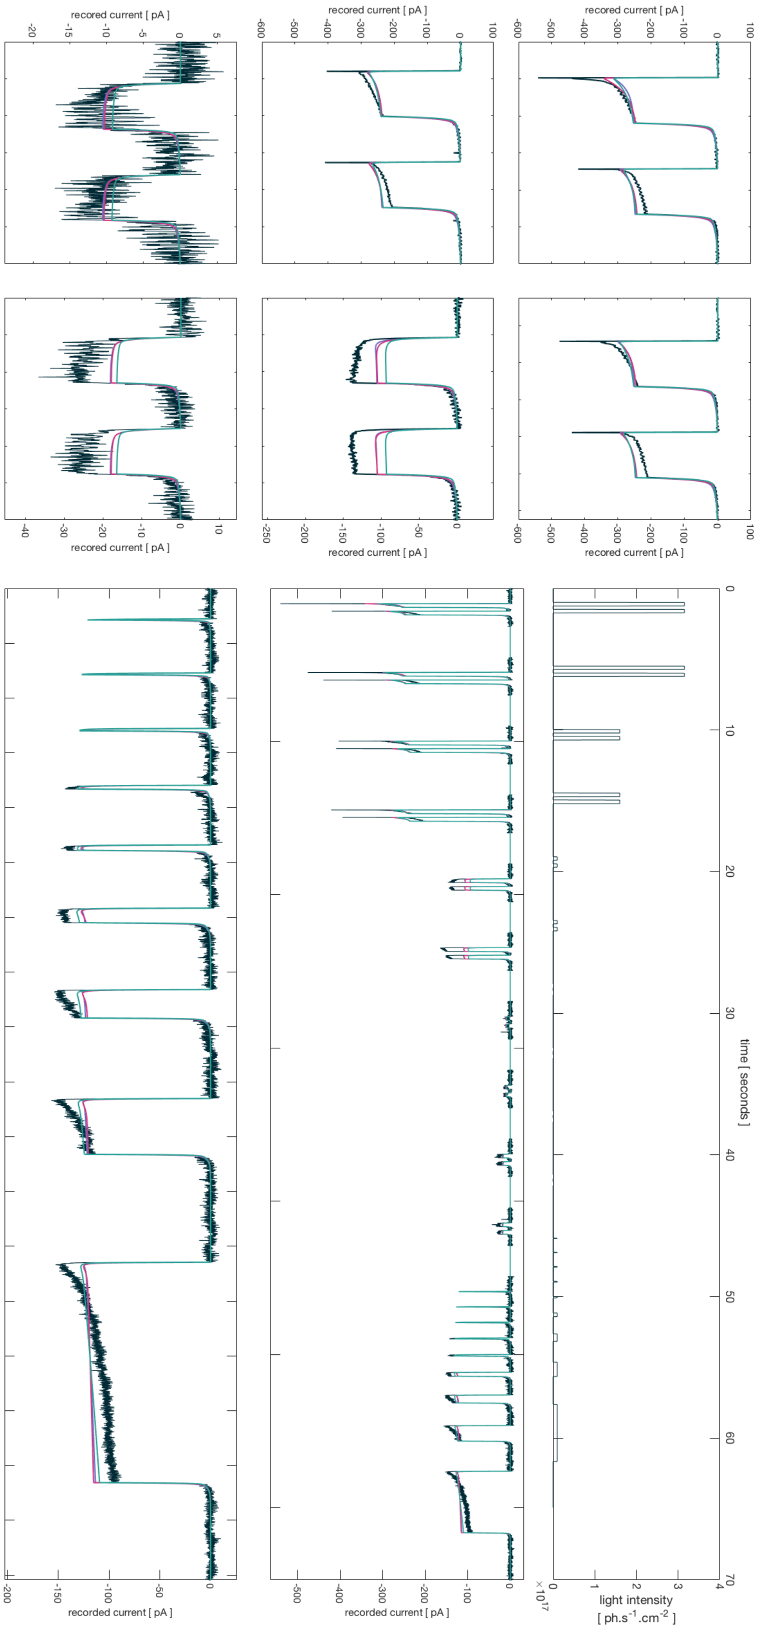

Figure 38: Validation on Cell #6

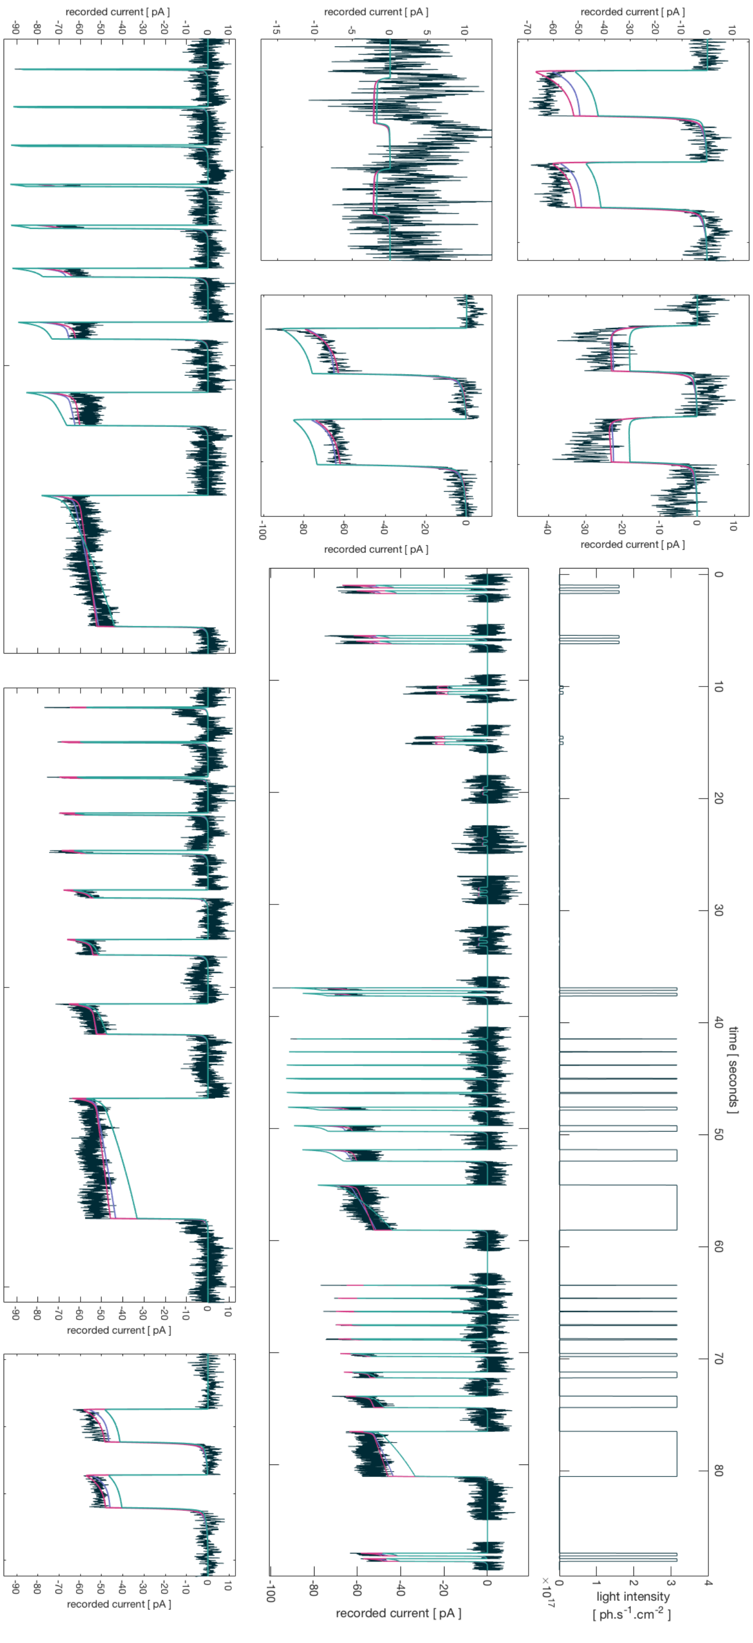

Figure 39: Validation on Cell #8

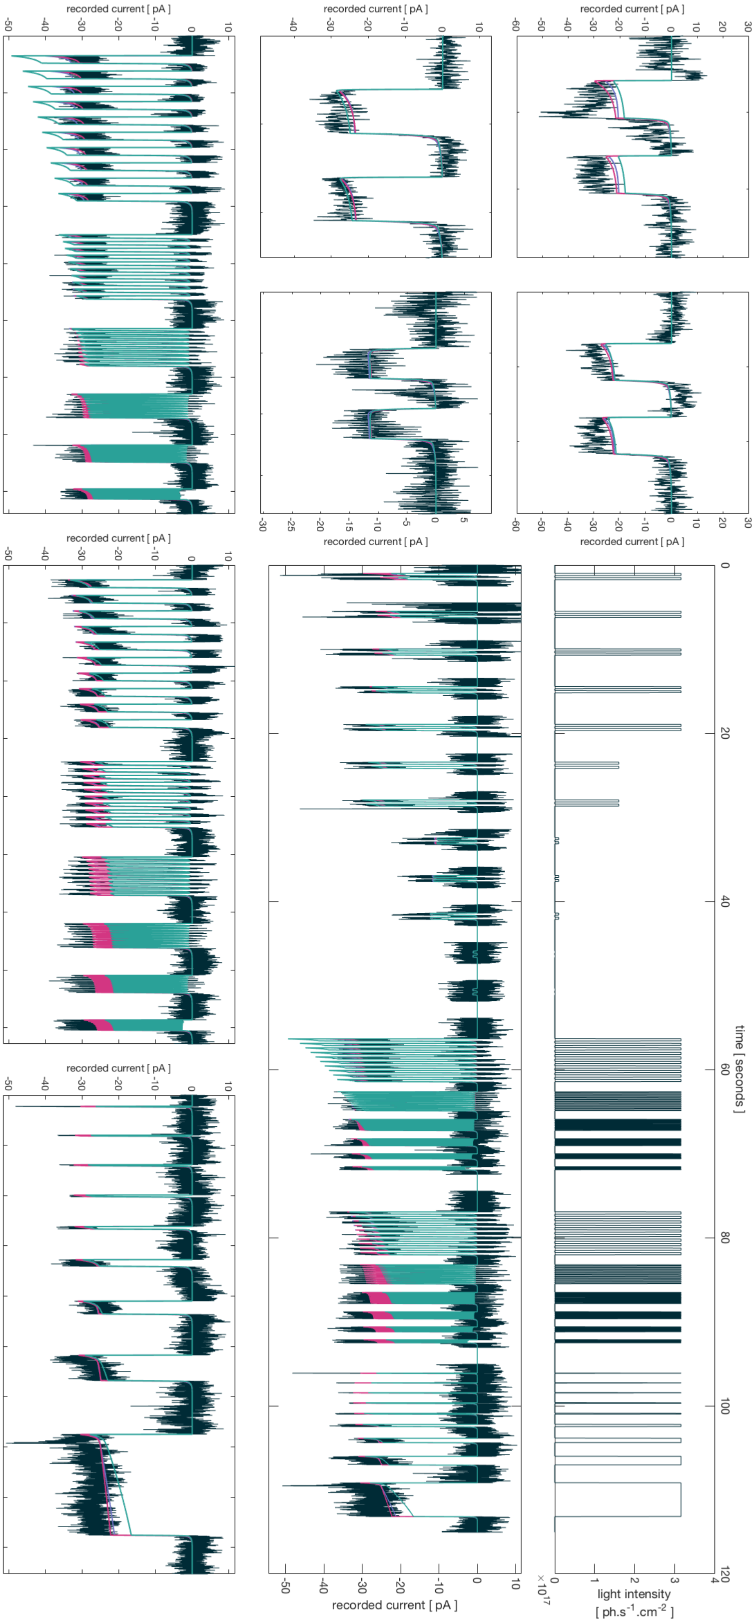

Figure 40: Validation on Cell #9

BIBLIOGRAPHY

[1] Powerfft asic. http://www.eonic.com/index.asp?item=32.

[2] Intel integrated performance primitives (ipp). http://software.intel.com/en-us/intel-ipp/, .

[3] Intel Math Kernel Library. http://software.intel.com/en-us/articles/intel-mkl/, .

[4] Bin Fang A, Yuefan Deng A, and Glenn Martyna B. Performance of the 3d fft on the 6d network torus qcdoc parallel supercomputer, 2007.

[5] Antoine R Adamantidis, Feng Zhang, Alexander M Aravanis, Karl Deisseroth, Luis De Lecea, and Luis de Lecea. Neural substrates of awakening probed with optogenetic control of hypocretin neurons. *Nature*, 450(7168):420–4, 2007. ISSN 1476-4687. doi: 10.1038/nature06310. URL `http://www.ncbi.nlm.nih.gov/pubmed/17943086`.

[6] Y. Ahmadian, A. M. Packer, R. Yuste, and L. Paninski. Designing optimal stimuli to control neuronal spike timing. *Journal of Neurophysiology*, 106(2):1038–1053, 2011. ISSN 0022-3077. doi: 10.1152/jn.00427.2010. URL `http://jn.physiology.org/cgi/doi/10.1152/jn.00427.2010`.

[7] Alexander M Aravanis, Li-Ping Wang, Feng Zhang, Leslie A Meltzer, Murtaza Z Mogri, M Bret Schneider, and Karl Deisseroth. An optical neural interface: in vivo control of rodent motor cortex with integrated fiberoptic and optogenetic technology. *Journal of Neural Engineering*, 4(3): S143, 2007. URL `http://stacks.iop.org/1741-2552/4/i=3/a=S02`.

[8] Christine Aurnhammer, Maren Haase, Nadine Muether, Martin Hausl, Christina Rauschhuber, Ingrid Huber, Hans Nitschko, Ulrich Busch, Andreas Sing, Anja Ehrhardt, and Armin Baiker. Universal real-time pcr for the detection and quantification of adeno-associated virus serotype 2-derived inverted terminal repeat sequences. *Human Gene Therapy Methods*, 23:18–28, 2012.

[9] B. Baas. A low-power, high-performance, 1024-point fft processor. In *IEEE Journal OF Solid-state Circuits*, volume 34, pages 380–387, 1999.

[10] Paul Bach-y Rita. Tactile sensory substitution studies. *Annals of the New York Academy of Sciences*, 1013:83–91, 2004. ISSN 0077-8923. doi: 10.1196/annals.1305.006.

[11] Paul Bach-y Rita and Stephen W. Kercel. Sensory substitution and the human-machine interface. *Trends in Cognitive Sciences*, 7(12):541–546, 2003. ISSN 13646613. doi: 10.1016/j.tics.2003.10.013.

[12] Paul Bach-y Rita, Carter C. Collins, Frank A. Saunders, Benjamin White, and Lawrence Scadden. Vision Substitution by Tactile Image Projection. *Nature*, 221:963–964, 1969. ISSN 0028-0836. doi: 10.1038/224488a0. URL `http://adsabs.harvard.edu/abs/1969Natur.224..177K`.

[13] Paul Bach-y Rita, Kurt A. Kaczmarek, Mitchell E. Tyler, and Jorge Garcia-Lara. Form perception with a 49-point electrotactile stimulus array on the tongue: A technical note. *J. Rehabil. Res. Dev.*, 35:427–430, 1998.

[14] Christian Bamann, Taryn Kirsch, Georg Nagel, and Ernst Bamberg. Spectral Characteristics of the Photocycle of Channelrhodopsin-2 and Its Implication for Channel Function. *Journal of Molecular Biology*, 375(3):686–694, 2008. ISSN 00222836. doi: 10.1016/j.jmb.2007.10.072.

[15] Christian Bamann, Ronnie Gueta, Sonja Kleinlogel, Georg Nagel, and Ernst Bamberg. Structural guidance of the photocycle of channelrhodopsin-2 by an interhelical hydrogen bond. *Biochemistry*, 49(2):267–278, 2010. ISSN 00062960. doi: 10.1021/bi901634p.

[16] R. Benosman, Sio-Hoï Ieng, P. Rogister, and C. Posch. Asynchronous event-based hebbian epipolar geometry. *Neural Networks, IEEE Transactions on*, 22(11):1723–1734, Nov 2011. ISSN 1045-9227. doi: 10.1109/TNN.2011.2167239.

[17] R. Benosman, C. Clercq, X. Lagorce, Sio-Hoi Ieng, and C. Bartolozzi. Event-based visual flow. *Neural Networks and*

Learning Systems, IEEE Transactions on, 25(2):407–417, Feb 2014. ISSN 2162-237X. doi: 10.1109/TNNLS.2013.2273537.

[18] André Berndt, Ofer Yizhar, Lisa A Gunaydin, Peter Hegemann, and Karl Deisseroth. Bi-stable neural state switches. *Nature neuroscience*, 12(2):229–34, 2009. ISSN 1546-1726. doi: 10.1038/nn.2247. URL http://dx.doi.org/10.1038/nn.2247.

[19] André Berndt, Matthias Prigge, Dietrich Gradmann, and Peter Hegemann. Two open states with progressive proton selectivities in the branched channelrhodopsin-2 photocycle. *Biophysical Journal*, 98(5):753–761, 2010. ISSN 00063495. doi: 10.1016/j.bpj.2009.10.052. URL http://dx.doi.org/10.1016/j.bpj.2009.10.052.

[20] R Berner, C Brandli, M Yang, and S C Liu. A 240-by-180 10mW 12us latency sparse-output vision sensor for mobile applications. In *VLSI Circuits (VLSIC), 2013 Symposium on*, pages 186–187, 2013. ISBN 9784863483484.

[21] M J Berry and M Meister. Refractoriness and neural precision. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 18(6):2200–11, 1998. ISSN 0270-6474. doi: 10.1016/S0315-5463(83)72188-7. URL http://www.ncbi.nlm.nih.gov/pubmed/9482804.

[22] E. L. Berson. Retinitis pigmentosa: The Friedenwald lecture. *Investigative Ophthalmology and Visual Science*, 34(5):1659–1676, 1993. ISSN 01460404.

[23] Edward S Boyden. Optogenetics and the future of neuroscience. *Nature Neuroscience*, 18(9):1200–1201, 2015. ISSN 1097-6256. doi: 10.1038/nn.4094. URL http://www.nature.com/doifinder/10.1038/nn.4094.

[24] Edward S Boyden, Feng Zhang, Ernst Bamberg, Georg Nagel, and Karl Deisseroth. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature neuroscience*, 8(9):1263–8, 2005. ISSN 1097-6256. doi: 10.1038/nn1525. URL http://www.ncbi.nlm.nih.gov/pubmed/16116447.

[25] Christian Brandli, Raphael Berner, Minhao Yang, Shih Chii Liu, and Tobi Delbruck. A 240-by-180 130dB 3us latency

global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits*, 49(10):2333–2341, 2014. ISSN 00189200. doi: 10.1109/JSSC.2014.2342715.

[26] Sara Bruun, Daniel Stoeppler, Anke Keidel, Uwe Kuhlmann, Meike Luck, Anne Diehl, Michel-andreas Geiger, David Woodmansee, Dirk Trauner, Peter Hegemann, Hartmut Oschkinat, Peter Hildebrandt, and Katja Stehfest. Light-Dark Adaptation of Channelrhodopsin Involves Photoconversion between the all- trans and 13-cis Retinal Isomers. *Biochemistry*, 2015. doi: 10.1021/acs.biochem.5b00597.

[27] G.T. Buracas, A.M. Zador, M.R. DeWeese, and T.D. Albright. Efficient dicrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron*, 20(5):959–969, 1998. ISSN 0896-6273. doi: 10.1016/S0896-6273(00)80477-8.

[28] Volker Busskamp, Jens Duebel, David Balya, Mathias Fradot, Tim James Viney, Sandra Siegert, Anna C Groner, Erik Cabuy, Valérie Forster, Mathias Seeliger, Martin Biel, Peter Humphries, Michel Paques, Saddek Mohand-Said, Didier Trono, Karl Deisseroth, José-Alain Sahel, Serge Picaud, and Botond Roska. Genetic reactivation of cone photoreceptors restores visual responses in retinitis pigmentosa. *Science*, 329(5990):413–7, July 2010. ISSN 1095-9203. doi: 10.1126/science.1190897. URL `http://www.ncbi.nlm.nih.gov/pubmed/20576849`.

[29] D. A. Butts, C. Weng, J. Jin, J.-M. Alonso, and L. Paninski. Temporal Precision in the Visual Pathway through the Interplay of Excitation and Stimulus-Driven Suppression. *Journal of Neuroscience*, 31(31):11313–11327, 2011. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.0434-11.2011. URL `http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.0434-11.2011`.

[30] Daniel A. Butts, Chong Weng, Jianzhong Jin, Chun I. Yeh, Nicholas A. Lesica, Jose Manuel Alonso, and Garrett B. Stanley. Temporal precision in the neural code and the timescales of natural vision. *Nature*, 449(7158):92–95, 2007. ISSN 14764687. doi: 10.1038/nature06105.

[31] Y. K. Chan and S. Y. Lim. Synthetic aperture radar (sar) signal generation. *Progress In Electromagnetics Research*, (B): 269–290, 2008.

[32] Thusitha N. Chandrapala and Bertram E. Shi. Invariant feature extraction from event based stimuli. *2016 6th IEEE International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, pages 1–6, 2016. doi: 10.1109/BIOROB.2016.7523449. URL `http://ieeexplore.ieee.org/document/7523449/`.

[33] Claudia Clopath, Tobias Bonhoeffer, Mark Hübener, and Tobias Rose. Variance and invariance of neuronal long-term representations. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1715):20160161, 2017. ISSN 0962-8436. doi: 10.1098/rstb.2016.0161. URL `http://rstb.royalsocietypublishing.org/lookup/doi/10.1098/rstb.2016.0161`.

[34] William T Cochran, James W Cooley, Davib L Favin, Howard D Helms, Reginald A Kaenel, William Lang, George C Maling, Davib E Nelson, Charles M Rader, and Peter D Welch. What is the Fast Fourier Transform. *IEEE transactions on audio and electroacoustics*, AU-15(2), 1967.

[35] J. Conradt, M. Cook, R. Berner, P. Lichtsteiner, R. J. Douglas, and T. Delbruck. A pencil balancing robot using a pair of AER dynamic vision sensors. In *Proceedings - IEEE International Symposium on Circuits and Systems*, pages 781–784, 2009. ISBN 9781424438280. doi: 10.1109/ISCAS.2009.5117867.

[36] Gabriel Cristobal, Laurent U. Perrinet, and Matthias S. Keil. Fundamentals. In *Biologically inspired Computer Vision: Fundamentals and Applications*, pages 1–6. Wiley-VCH Verlag GmBH & Co, 1 edition, 2016.

[37] Dennis M Dacey. The mosaic of midget ganglion cells in the human retina. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 13(12):5334–5355, 1993. ISSN 0270-6474.

[38] Paolo D’Alberto, Peter A. Milder, Aliaksei Sandryhaila, Franz Franchetti, James C. Hoe, Jose M.F. Moura, Markus

Puschel, and Jeremy R. Johnson. Generating fpga-accelerated dft libraries. *Field-Programmable Custom Computing Machines, Annual IEEE Symposium on*, 0:173–184, 2007. doi: http://doi.ieeecomputersociety.org/10.1109/FCCM.2007.58.

[39] D. Dalkara, L. C. Byrne, R. R. Klimczak, M. Visel, L. Yin, W. H. Merigan, J. G. Flannery, and D. V. Schaffer. In Vivo-Directed Evolution of a New Adeno-Associated Virus for Therapeutic Outer Retinal Gene Delivery from the Vitreous. *Science Translational Medicine*, 5(189):189ra76, 2013. ISSN 1946-6234. doi: 10.1126/scitranslmed.3005708. URL `http://stm.sciencemag.org/cgi/doi/10.1126/scitranslmed.3005708`.

[40] Karl Deisseroth. Optogenetics : 10 years of microbial opsins in neuroscience. *Nature neuroscience*, 18(9):1213–1225, 2015.

[41] Tobi Delbruck. Neuromorophic vision sensing and processing. *European Solid-State Device Research Conference*, 2016-Octob:7–14, 2016. ISSN 19308876. doi: 10.1109/ESSDERC.2016.7599576.

[42] Alain Destexhe, Zf F Mainen, and Tj J Sejnowski. Synthesis of models for excitable membranes, synaptic transmission and neuromodulation using a common kinetic formalism. *Journal of computational neuroscience*, 1(3):195–230, 1994. ISSN 0929-5313. doi: 10.1007/BF00961734. URL `http://link.springer.com/article/10.1007/BF00961734`.

[43] Giulia V. Elli, Stefania Benetti, and Olivier Collignon. Is there a future for sensory substitution outside academic laboratories? *Multisensory Research*, 27(5-6):271–291, 2014. ISSN 22134808. doi: 10.1163/22134808-00002460.

[44] Flaxman *et. al.*. Global causes of blindness and distance vision impairment 1990-2020: a systematic review and meta-analysis. *The Lancet Global Health*, 5(12):e1221–e1234, 2017. ISSN 2214109X. doi: 10.1016/S2214-109X(17)30393-5.

[45] Frida Eng. *Non-Uniform Sampling in Statistical Signal Processing*. PhD thesis, 2007.

[46] Frida Eng and Fredrik Gustafsson. Frequency Transforms based on Nonuniform Sampling - Basic Stochastic Properties. *Proceedings of Radiovetenskap och kommunikation 2005*, pages 1–4, 2005.

[47] Benjamin D Evans, Sarah Jarvis, Simon R Schultz, and Konstantin Nikolic. PyRhO : A Multiscale Optogenetics Simulation Platform. *Frontiers in Neuroinformatics*, 10 (March):1–19, 2016. doi: 10.3389/fninf.2016.00008.

[48] Mohsen Firouzi and Jörg Conradt. Asynchronous Event-based Cooperative Stereo Matching Using Neuromorphic Silicon Retinas. *Neural Processing Letters*, 43(2):311–326, 2016. ISSN 1573773X. doi: 10.1007/s11063-015-9434-5.

[49] Franz Franchetti, Markus Püschel, Yevgen Voronenko, Srinivas Chellappa, and José M. F. Moura. Discrete fourier transform on multicore. *IEEE Signal Processing Magazine, special issue on Signal Processing on Platforms with Multiple Cores*, 26(6):90–102, 2009.

[50] Yves Frégnac. Neurobiology - Homeostasis or synaptic plasticity? *Nature*, 391(6670):845–846, 1998. ISSN 00280836. doi: 10.1038/35996.

[51] Yves Frégnac. Big data and the industrialization of neuroscience: A safe roadmap for understanding the brain? *Science*, 358(6362):470–477, 2017. ISSN 10959203. doi: 10.1126/science.aan8866.

[52] Matteo Frigo and Steven G. Johnson. FFTW: An adaptive software architecture for the FFT. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, volume 3, pages 1381–1384, Seattle, Washington, 1998.

[53] Steve Furber. Large-scale neuromorphic computing systems. *Journal of Neural Engineering*, 13(5), 2016. ISSN 17412552. doi: 10.1088/1741-2560/13/5/051001.

[54] Steve B. Furber, Francesco Galluppi, Steve Temple, and Luis a. Plana. The SpiNNaker Project. *Proceedings of the IEEE*, 102(5):652–665, May 2014. ISSN 0018-9219. doi: 10.1109/JPROC.2014.2304638. URL `http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6750072`.

[55] Francesco Galluppi, Didier Pruneau, Joel Chavas, Xavier Lagorce, Christoph Posch, Guillaume Chenegros, Gilles Cordurie, Charlie Galle, Nicolas Oddo, and Ryad Benosman. A stimulation platform for optogenetic and bionic vision restoration. In *Proceedings - IEEE International Symposium on Circuits and Systems*, pages 3–6, 2017. ISBN 9781467368520. doi: 10.1109/ISCAS.2017.8050683.

[56] Steven J Gray, Vivian W Choi, Aravind Asokan, Rebecca A Haberman, J Thomas, and Richard Jude Samulski. *Production of Recombinant Adeno-Associated Viral Vectors and Use in In Vitro and In Vivo Administration*. 2011. ISBN 0471142301. doi: 10.1002/0471142301.ns0417s57.Production.

[57] Modris Greitans and Rolands Shavelis. Extended Fourier series for time-varying filtering and reconstruction from level-crossing samples. *European Signal Processing Conference*, (4):1–5, 2013. ISSN 22195491.

[58] Nir Grossman, Konstantin Nikolic, Christofer Toumazou, and Patrick Degenaar. Modeling study of the light stimulation of a neuron cell with channelrhodopsin-2 mutants. *IEEE Transactions on Biomedical Engineering*, 58(6): 1742–1751, 2011. ISSN 00189294. doi: 10.1109/TBME.2011.2114883.

[59] Lisa A. Gunaydin, Ofer Yizhar, André Berndt, Vikaas S. Sohal, Karl Deisseroth, and Peter Hegemann. Ultrafast optogenetic control. *Nature Neuroscience*, 13(3):387–392, 2010. ISSN 10976256. doi: 10.1038/nn.2495. URL `http://dx.doi.org/10.1038/nn.2495`.

[60] Dyonne T. Hartong, Eliot L. Berson, and Thaddeus P. Dryja. Retinitis pigmentosa. *Lancet*, 368(9549):1795–1809, 2006. ISSN 01406736. doi: 10.1016/S0140-6736(06)69740-7.

[61] T. Hawkes and P. Simonpieri. Signal coding using asynchronous delta modulation. *IEEE Trans. Commun*, (5):729–731, 1974.

[62] Peter Hegemann. Vision in microalgae. *Planta*, 203(3):265–274, 1997. ISSN 00320935. doi: 10.1007/s004250050191.

[63] Peter Hegemann, Sabine Ehlenbeck, and Dietrich Gradmann. Multiple photocycles of channelrhodopsin.

Biophysical journal, 89(6):3911–8, 2005. ISSN 0006-3495. doi: 10.1529/biophysj.105.069716. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1366958{&}tool=pmcentrez{&}rendertype=abstract.

[64] B. Heimler, N. Weisz, and O. Collignon. Revisiting the adaptive and maladaptive effects of crossmodal plasticity. *Neuroscience*, 283:44–63, 2014. ISSN 18737544. doi: 10.1016/j.neuroscience.2014.08.003. URL http://dx.doi.org/10.1016/j.neuroscience.2014.08.003.

[65] Shoko Hososhima, Seiichiro Sakai, Toru Ishizuka, and Hiromu Yawo. Kinetic evaluation of photosensitivity in Bistable variants of chimeric channelrhodopsins. *PLoS ONE*, 10(3):1–14, 2015. ISSN 19326203. doi: 10.1371/journal.pone.0119558.

[66] Mark S Humayun, Jessy D Dorn, Lyndon Cruz, and Gislin Dagnelie. Interim Results from the International Trial of Second Sight ' s Visual Prosthesis. *American Academy of Ophtalmology*, 119(4):779–788, 2012. doi: 10.1016/j.ophtha.2011.09.028.

[67] Sio-Hoi Ieng, Christoph Posch, and Ryad Benosman. Asynchronous Neuromorphic Event-Driven Image Filtering. *Proceedings of the IEEE*, 102(10), 2014. ISSN 0018-9219. doi: 10.1109/JPROC.2014.2347355.

[68] Giacomo Indiveri and Shih Chii Liu. Memory and Information Processing in Neuromorphic Systems. *Proceedings of the IEEE*, 103(8):1379–1397, 2015. ISSN 00189219. doi: 10.1109/JPROC.2015.2444094.

[69] L. H. Jepson, 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. *Journal of Neuroscience*, 33(17):7194–7205, 2013. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.4967-12.2013. URL http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.4967-12.2013.

[70] L. H. Jepson, P. Hottowy, K. Mathieson, D. E. Gunning, W. DÂ browski, A. M. Litke, and E. J. Chichilnisky. Spa-

tially Patterned Electrical Stimulation to Enhance Resolution of Retinal Prostheses. *Journal of Neuroscience*, 34 (14):4871–4881, 2014. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.2882-13.2014. URL `http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.2882-13.2014`.

[71] Lauren H. Jepson, Pawel Hottowy, Geoffrey A. Weiner, Władysław Dabrowski, Alan M. Litke, and E. J. Chichilnisky. High-fidelity reproduction of spatiotemporal visual signals for retinal prosthesis. *Neuron*, 83(1):87–92, 2014. ISSN 10974199. doi: 10.1016/j.neuron.2014.04.044.

[72] Renaud Jolivet, Timothy J Lewis, and Wulfram Gerstner. Generalized integrate-and-fire models of neuronal activity approximate spike trains of a detailed model to a high degree of accuracy. *Journal of neurophysiology*, 92(2):959–976, 2004. ISSN 0022-3077. doi: 10.1152/jn.00190.2004.

[73] Bryan W Jones and Robert E Marc. Retinal remodeling during retinal degeneration. *Experimental eye research*, 81 (2):123–37, aug 2005. ISSN 0014-4835. doi: 10.1016/j.exer.2005.03.006. URL `http://www.ncbi.nlm.nih.gov/pubmed/15916760`.

[74] Hideaki E Kato, Feng Zhang, Ofer Yizhar, Charu Ramakrishnan, Tomohiro Nishizawa, Kunio Hirata, Jumpei Ito, Karl Deisseroth, and Osamu Nureki. Crystal structure of the channelrhodopsin light-gated cation channel. *Nature*, 482:369–374, 2012. doi: 10.1038/nature10870.

[75] N Katsanis, Sj Ansley, Jl Badano, Er Eichers, Ra Lewis, Pj Scambler, Ws Davidson, Pl Beales, and Lupski. Evidence for triallelic inheritance in Bardet-Biedl syndrome, a recessive, genetically heterogeneous and pleiotropic disorder. 293(SEPTEMBER):2256–2259, 2001. URL `http://discovery.ucl.ac.uk/88111/`.

[76] Justin Keat, Pamela Reinagel, R. Clay Reid, and Markus Meister. Predicting every spike: A model for the responses of visual neurons. *Neuron*, 30(3):803–817, 2001. ISSN 08966273. doi: 10.1016/S0896-6273(01)00322-1.

[77] Sihem Kime, Francesco Galluppi, Xavier Lagorce, Ryad B. Benosman, and Jean Lorenceau. Psychophysical Assessment of Perceptual Performance with Varying Display

Frame Rates. *Journal of Display Technology*, 12(11):1372–1382, 2016. ISSN 1551319X. doi: 10.1109/JDT.2016.2603222.

[78] Nathan C Klapoetke, Yasunobu Murata, Sung Soo Kim, Stefan R Pulver, Amanda Birdsey-Benson, Yong Ku Cho, Tania K Morimoto, Amy S Chuong, Eric J Carpenter, Zhijian Tian, Jun Wang, Yinlong Xie, Zhixiang Yan, Yong Zhang, Brian Y Chow, Barbara Surek, Michael Melkonian, Vivek Jayaraman, Martha Constantine-Paton, Gane Ka-Shu Wong, and Edward S Boyden. Independent optical excitation of distinct neural populations. *Nature Methods*, 11(3):338–346, 2014. ISSN 1548-7091. doi: 10.1038/nmeth.2836. URL `http://www.nature.com/doifinder/10.1038/nmeth.2836`.

[79] Jürgen Kogler, Christoph Sulzbachner, Florian Eibensteiner, and Martin Humenberger. Address-Event Matching for a Silicon Retina Based Stereo Vision System. In *International Conference from Scientific Computing to Computational Engineering*, 2010. ISBN 9789533075167.

[80] Benjamin S. Krause, Christiane Grimm, Joel C.D. Kaufmann, Franziska Schneider, Thomas P. Sakmar, Franz J. Bartl, and Peter Hegemann. Complex Photochemistry within the Green-Absorbing Channelrhodopsin ReaChR. *Biophysical Journal*, 112(6):1166–1175, 2017. ISSN 15420086. doi: 10.1016/j.bpj.2017.02.001. URL `http://dx.doi.org/10.1016/j.bpj.2017.02.001`.

[81] Jens Kuhne, Kirstin Eisenhauer, Eglof Ritter, Peter Hegemann, Klaus Gerwert, and Franz Bartl. Early Formation of the Ion-Conducting Pore in Channelrhodopsin-2. *Angewandte Chemie*, 54:4953–4957, 2015. doi: 10.1002/anie.201410180.

[82] R. D. Kumbhani, M. J. Nolt, and L. A. Palmer. Precision, Reliability, and Information-Theoretic Analysis of Visual Thalamocortical Neurons. *Journal of Neurophysiology*, 98(5): 2647–2663, 2007. ISSN 0022-3077. doi: 10.1152/jn.00900.2006. URL `http://jn.physiology.org/cgi/doi/10.1152/jn.00900.2006`.

[83] Pinit Kumhom. *Design, Optimization, and Implementation of a Universal FFT Processor*. PhD thesis, Electrical and Com-

puter Engineering, Drexel University, 2001. Also Tech. Report DU-MCS-01-01, Drexel University, 2001.

[84] X. Lagorce, C. Meyer, S.-H. Ieng, D. Filliat, and R. Benosman. Asynchronous event-based multikernel algorithm for high-speed visual features tracking. *Neural Networks and Learning Systems, IEEE Transactions on*, PP(99):1–1, 2014. ISSN 2162-237X. doi: 10.1109/TNNLS.2014.2352401.

[85] Xavier Lagorce, Cédric Meyer, Sio-Hoi Ieng, David Filliat, and Ryad Benosman. Asynchronous Event-Based Multikernel Algorithm for High Speed Visual Feature Tracking. *IEEE transactions on neural networks and learning systems*, 26 (8):1710–1720, 2015. ISSN 2162-2388. doi: 10.1109/TNNLS.2014.2352401.

[86] T. Lenart, M. Gustafsson, and V. Owall. A hardware acceleration platform for digital holographic imaging. *Journal of Signal Processing System*, 52(3):297–311, 2008.

[87] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. A 128x128 120dB 15 us latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits*, 43(2):566–576, 2008.

[88] Bin Lin, Richard H Masland, and Enrica Strettoi. Remodeling of cone photoreceptor cells after rod degeneration in rd mice. *Experimental eye research*, 88(3):589–599, 2009. doi: 10.1016/j.exer.2008.11.022.

[89] Yu-Wei Lin, Hsuan-Yu Liu, and Chen-Yi Lee. A 1-gs/s fft/ifft processor for uwb applications. In *IEEE Journal of Solid-State Circuits*, pages 1726–1735, 2005.

[90] Robert C Liu, Svilen Tzonev, Sergei Rebrik, Kenneth D Miller, C Robert, Svilen Tzonev, Sergei Rebrik, and D Kenneth. Variability and Information in a Neural Code of the Cat Lateral Geniculate Nucleus. *Journal of Neurophysiology*, 86:2789–2806, 2001.

[91] Henri Lorach, Ryad Benosman, Olivier Marre, Sio-Hoi Ieng, JosÃ© A Sahel, and Serge Picaud. Artificial retina: the multichannel processing of the mammalian retina achieved with a neuromorphic asynchronous light acquisition device. *Journal of Neural Engineering*, 9(6):066004,

2012. URL http://stacks.iop.org/1741-2552/9/i=6/a=066004.

[92] Henri Lorach, Olivier Marre, JosÃ©-Alain Sahel, Ryad Benosman, and Serge Picaud. Neural stimulation for visual rehabilitation: Advances and challenges. *Journal of Physiology-Paris*, 107(5):421 – 431, 2013. ISSN 0928-4257. doi: http://dx.doi.org/10.1016/j.jphysparis.2012.10.003. URL http://www.sciencedirect.com/science/article/pii/S0928425712000678. Special issue: Neural Coding and Natural Image Statistics.

[93] V. A. Lorenz-Fonfria, T. Resler, N. Krause, M. Nack, M. Gossing, G. Fischer von Mollard, C. Bamann, E. Bamberg, R. Schlesinger, and J. Heberle. Transient protonation changes in channelrhodopsin-2 and their relevance to channel gating. *Proceedings of the National Academy of Sciences*, 110(14):E1273–E1281, 2013. ISSN 0027-8424. doi: 10.1073/pnas.1219502110. URL http://www.pnas.org/cgi/doi/10.1073/pnas.1219502110.

[94] Víctor A. Lórenz-Fonfría and Joachim Heberle. Channelrhodopsin unchained: Structure and mechanism of a light-gated cation channel. *Biochimica et Biophysica Acta - Bioenergetics*, 1837(5):626–642, 2014. ISSN 18792650. doi: 10.1016/j.bbabio.2013.10.014. URL http://dx.doi.org/10.1016/j.bbabio.2013.10.014.

[95] Shachar Maidenbaum, Sami Abboud, and Amir Amedi. Sensory substitution: Closing the gap between basic research and widespread practical visual rehabilitation. *Neuroscience and Biobehavioral Reviews*, 41:3–15, 2014. ISSN 18737528. doi: 10.1016/j.neubiorev.2013.11.007. URL http://dx.doi.org/10.1016/j.neubiorev.2013.11.007.

[96] Robert E Marc, Bryan W Jones, Carl B Watt, and Enrica Strettoi. Neural remodeling in retinal degeneration. *Progress in Retinal and Eye Research*, 22(5):607–655, sep 2003. ISSN 13509462. doi: 10.1016/S1350-9462(03)00039-9. URL http://linkinghub.elsevier.com/retrieve/pii/S1350946203000399.

[97] Carver Mead. Neuromorphic Electronic Systems. *Proceedings of the IEEE*, 78(10):1629–1636, 1990. ISSN 15582256. doi: 10.1109/5.58356.

[98] Ann H. Milam, Zong-Yi Li, and Robert N. Fariss. Histopathology of the human retina in retinitis pigmentosa. *Progress in retinal and eye research*, 17(2):175–205, April 1998. ISSN 1350-9462. URL `http://www.ncbi.nlm.nih.gov/pubmed/9695792`.

[99] Peter A. Milder, Franz Franchetti, James C. Hoe, and Markus Püschel. Formal datapath representation and manipulation for implementing dsp transforms. In *Proceedings of the 45th Annual Design Automation Conference*, DAC '08, pages 385–390, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-115-6. doi: 10.1145/1391469.1391572. URL `http://doi.acm.org/10.1145/1391469.1391572`.

[100] Maria Müller, Christian Bamann, Ernst Bamberg, and Werner Kühlbrandt. Projection structure of channelrhodopsin-2 at 6 Å resolution by electron crystallography. *Journal of Molecular Biology*, 414(1):86–95, 2011. ISSN 00222836. doi: 10.1016/j.jmb.2011.09.049.

[101] Georg Nagel, Doris Ollig, Markus Fuhrmann, Suneel Kateriya, Anna Maria Musti, Ernst Bamberg, and Peter Hegemann. Channelrhodopsin-1 : A Light-Gated Proton Channel in Green Algae. *Science*, 296(5577):2395–2398, 2002.

[102] Georg Nagel, Tanjef Szellas, Wolfram Huhn, Suneel Kateriya, Nona Adeishvili, Peter Berthold, Doris Ollig, Peter Hegemann, and Ernst Bamberg. Channelrhodopsin-2, a directly light-gated cation-selective membrane channel. *Proceedings of the National Academy of Sciences of the United States of America*, 100(24):13940–5, 2003. ISSN 0027-8424. doi: 10.1073/pnas.1936192100. URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=283525{&}tool=pmcentrez{&}rendertype=abstract`.

[103] Zhenjiang Ni, A. Bolopion, J. Agnus, R. Benosman, and S. Regnier. Asynchronous event-based visual shape tracking for stable haptic feedback in microrobotics. *Robotics, IEEE Transactions on*, 28(5):1081–1089, Oct 2012. ISSN 1552-3098. doi: 10.1109/TRO.2012.2198930.

[104] Zhenjiang Ni, C. Pacoret, Ryad Benosman, S.H. Ieng, and Stéphane Régnier. Asynchronous event-based high

speed vision for microparticle tracking. *Journal of Microscopy*, 245(3):236–244, 2012. ISSN 1365-2818. doi: 10.1111/j.1365-2818.2011.03565.x. URL `http://dx.doi.org/10.1111/j.1365-2818.2011.03565.x`.

[105] Konstantin Nikolic, Nir Grossman, Matthew S. Grubb, Juan Burrone, Chris Toumazou, and Patrick Degenaar. Photocycles of channelrhodopsin-2. *Photochemistry and Photobiology*, 85(1):400–411, 2009. ISSN 00318655. doi: 10.1111/j.1751-1097.2008.00460.x.

[106] J. M. Ong and L. Da Cruz. A review and update on the current status of stem cell therapy and the retina. *British Medical Bulletin*, 102(1):133–146, 2012. ISSN 00071420. doi: 10.1093/bmb/lds013.

[107] G. Orchard, C. Meyer, R. Etienne-Cummings, C. Posch, N. Thakor, and R. Benosman. Hfirst: A temporal approach to object recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PP(99):1–1, 2015. ISSN 0162-8828. doi: 10.1109/TPAMI.2015.2392947.

[108] L Paninski, J W Pillow, and E P Simoncelli. Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model. *Neural Comput*, 16(12):2533–2561, 2004. ISSN 0899-7667. doi: 10.1162/0899766042321797.

[109] C. L. Passaglia and John B. Troy. Information Transmission Rates of Cat Retinal Ganglion Cells. *Journal of Neurophysiology*, 91(3):1217–1229, 2003. ISSN 0022-3077. doi: 10.1152/jn.00796.2003. URL `http://jn.physiology.org/cgi/doi/10.1152/jn.00796.2003`.

[110] Xi Peng, Bo Zhao, Rui Yan, Huajin Tang, and Zhang Yi. Bag of events: An efficient probability-based feature extraction method for AER image sensors. *IEEE Transactions on Neural Networks and Learning Systems*, 28(4):791–803, 2017. ISSN 21622388. doi: 10.1109/TNNLS.2016.2536741.

[111] Laurent U. Perrinet. Role of homeostasis in learning sparse representations. *Neural Computation*, 22(7):1812–1836, 2010. doi: 10.1162/neco.2010.05-08-795. URL `https://doi.org/10.1162/neco.2010.05-08-795`. PMID: 20235818.

[112] Niclas Persson. *Event Based Sampling with Application to Spectral Estimation*. PhD thesis, 2007.

[113] J. W. Pillow, Liam Paninski, Valerie J. Uzzell, Eero P. Simoncelli, and E. J. Chichilnisky. Prediction and Decoding of Retinal Ganglion Cell Responses with a Probabilistic Spiking Model. *Journal of Neuroscience*, 25(47):11003–11013, 2005. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.3305-05.2005. URL `http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.3305-05.2005`.

[114] Jonathan W. Pillow, Yashar Ahmadian, and Liam Paninski. Model-Based Decoding, Information Estimation, and Change-Point Detection Techniques for Multineuron Spike Trains. *Neural computation*, 45(1):1–45, 2010. ISSN 1530-888X. doi: 10.1162/NECO_a_00058. URL `http://www.ncbi.nlm.nih.gov/pubmed/20964538{%}5Cnhttp://dx.doi.org/10.1162/NECO{_}a{_}00058{%}5Cnhttp://www.mitpressjournals.org/doi/abs/10.1162/NECO{_}a{_}00058{#}.VQgNMRCsVCw{%}5Cnhttp://www.mitpressjournals.org/doi/pdf/10.1162/NECO{_}a{_}00058{%}23.VQgNMRCsVCw`.

[115] Christoph Posch, Daniel Matolin, and Rainer Wohlgenannt. An asynchronous time-based image sensor. *2008 IEEE International Symposium on Circuits and Systems*, pages 2130–2133, May 2008. doi: 10.1109/ISCAS.2008.4541871. URL `http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4541871`.

[116] Maurice Ptito, Isabelle Matteau, Albert Gjedde, and Ron Kupers. Recruitment of the middle temporal area by tactile motion in congenital blindness. *NeuroReport*, 20: 543–547, 2009. ISSN 0959-4965. doi: 10.1097/WNR.0b013e3283279909.

[117] Markus Püschel, José M. F. Moura, Jeremy Johnson, David Padua, Manuela Veloso, Bryan Singer, Jianxin Xiong, Franz Franchetti, Aca Gacic, Yevgen Voronenko, Kang Chen, Robert W. Johnson, and Nicholas Rizzolo. Spiral: Code generation for DSP transforms. *Proceedings of the IEEE, special issue on Program Generation, Optimization, and Adaptation*, 93(2):232–275, 2005.

[118] Zhuo Qian and Martin Margala. A novel low-power and in-place split-radix fft processor. In *Proceedings of the 24th Edition of the Great Lakes Symposium on VLSI*, GLSVLSI '14, pages 81–82, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2816-6. doi: 10.1145/2591513.2591563. URL http://doi.acm.org/10.1145/2591513.2591563.

[119] Ionela Radu, Christian Bamann, Melanie Nack, Georg Nagel, Ernst Bamberg, and Joachim Heberle. Conformational Changes of Channelrhodopsin-2. *The Journal of the American Chemical Society*, 446(7136):7313–7319, 2009.

[120] P Reinagel and R C Reid. Temporal coding of visual information in the thalamus. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 20(14):5392–400, 2000. ISSN 0270-6474. doi: 20/14/5392[pii]. URL http://www.ncbi.nlm.nih.gov/pubmed/10884324.

[121] Eglof Ritter, Patrick Piwowarski, Peter Hegemann, and Franz J Bartl. Light-dark Adaptation of Channelrhodopsin C128T Mutant. *The Journal of Biological Chemistry*, 288(15): 10451–10458, 2013. doi: 10.1074/jbc.M112.446427.

[122] P. Rogister, R. Benosman, Sio-Hoi Ieng, P. Lichtsteiner, and T. Delbruck. Asynchronous event-based binocular stereo matching. *Neural Networks and Learning Systems, IEEE Transactions on*, 23(2):347–353, Feb 2012. ISSN 2162-237X. doi: 10.1109/TNNLS.2011.2180025.

[123] Bodo Rueckauer and Tobi Delbruck. Evaluation of event-based algorithms for optical flow with ground-truth from inertial measurement sensor. *Frontiers in Neuroscience*, 10 (APR):1–17, 2016. ISSN 1662453X. doi: 10.3389/fnins.2016.00176.

[124] José-Alain Sahel and Botond Roska. Gene Therapy for Blindness. *Annual Review of Neuroscience*, 36(1):467–488, 2013. ISSN 0147-006X. doi: 10.1146/annurev-neuro-062012-170304. URL http://www.annualreviews.org/doi/10.1146/annurev-neuro-062012-170304.

[125] Conrad Sanderson and Ryan Curtin. Armadillo: a template-based C++ library for linear algebra. *The Jour-*

nal of Open Source Software, 1:26, 2016. doi: 10.21105/joss.00026.

[126] Joshua R. Sanes and Richard H. Masland. The Types of Retinal Ganglion Cells: Current Status and Implications for Neuronal Classification. *Annual Review of Neuroscience*, 38(1):221–246, 2015. ISSN 0147-006X. doi: 10.1146/annurev-neuro-071714-034120. URL `http://www.annualreviews.org/doi/10.1146/annurev-neuro-071714-034120`.

[127] Franziska Schneider, Dietrich Gradmann, and Peter Hegemann. Ion selectivity and competition in channelrhodopsins. *Biophysical Journal*, 105(1):91–100, 2013. ISSN 00063495. doi: 10.1016/j.bpj.2013.05.042. URL `http://dx.doi.org/10.1016/j.bpj.2013.05.042`.

[128] Franziska Schneider, Christiane Grimm, and Peter Hegemann. Biophysics of Channelrhodopsin. *Annual Review of Biophysics*, 2015. doi: 10.1146/annurev-biophys-060414-034014.

[129] M. S. Singh and R. E. MacLaren. Stem cells as a therapeutic tool for the blind: biology and future prospects. *Proceedings of the Royal Society B: Biological Sciences*, 278(1721): 3009–3016, 2011. ISSN 0962-8452. doi: 10.1098/rspb.2011.1028. URL `http://rspb.royalsocietypublishing.org/cgi/doi/10.1098/rspb.2011.1028`.

[130] Katja Stehfest and Peter Hegemann. Evolution of the Channelrhodopsin Photocycle Model. *ChemPhysChem*, pages 1120–1126, 2010. doi: 10.1002/cphc.200900980.

[131] a Szucs. Applications of the spike density function in analysis of neuronal firing patterns. *Journal of neuroscience methods*, 81(1-2):159–167, 1998. ISSN 0165-0270. doi: 10.1016/S0165-0270(98)00033-8.

[132] Dillon T. Two virtex-ii fpgas deliver fastest, cheapest, best high-performance image processing system. In *Xilinx Xcell, Journal*, volume 41, pages 70–73, 2001.

[133] Yannis Tsividis. Event-Driven Data Acquisition and Digital Signal Processing - A Tutorial. *IEEE Transactions on*

Circuits and Systems II: Express Briefs, 57(8):577–581, 2010. ISSN 1549-7747. doi: 10.1109/TCSII.2010.2056012.

[134] I. Uzun, A. Amira, and A. Bouridane. Fpga implementations of fast fourier transforms for real-time signal and image processing. In *IEEE Proceedings. Vision, Image, and Signal Processing*, volume 152, pages 283–296, 2005.

[135] V. J. Uzzell and E. J. Chichilnisky. Precision of Spike Trains in Primate Retinal Ganglion Cells. *Journal of Neurophysiology*, 92(2):780–789, 2004. ISSN 0022-3077. doi: 10.1152/jn.01171.2003. URL http://jn.physiology.org/cgi/doi/10.1152/jn.01171.2003.

[136] Charles Van Loan. *Computational frameworks for the fast Fourier transform*. SIAM, Philadelphia, 1992.

[137] C. Vezyrtzis and Y. Tsividis. Processing of signals using level-crossing sampling. *Proc. IEEE Int. Symp. Circuits Syst*, (1):2293–2296, 2009.

[138] J D Victor. Temporal aspects of neural coding in the retina and lateral geniculate. *Network: Computation in Neural Systems*, 10(4), 1999. ISSN 0954-898X. doi: 10.1088/0954-898X/10/4/201.

[139] Zhou Wang, Alan Conrad Bovik, Hamid Rahim Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, 13(4):600–12, April 2004. ISSN 1057-7149. URL http://www.ncbi.nlm.nih.gov/pubmed/15376593.

[140] Jamie Ward and Peter Meijer. Visual experiences in the blind induced by an auditory sensory substitution device. *Consciousness and Cognition*, 19(1):492–500, 2010. ISSN 10538100. doi: 10.1016/j.concog.2009.10.006. URL http://dx.doi.org/10.1016/j.concog.2009.10.006.

[141] Hiroshi C. Watanabe, Kai Welke, Daniel J. Sindhikara, Peter Hegemann, and Marcus Elstner. Towards an understanding of channelrhodopsin function: Simulations lead to novel insights of the channel mechanism. *Journal of Molecular Biology*, 425(10):1795–1814, 2013. ISSN 00222836.

doi: 10.1016/j.jmb.2013.01.033. URL `http://dx.doi.org/10.1016/j.jmb.2013.01.033`.

[142] Zhen Xie, Shengyong Chen, and Garrick Orchard. Event-based stereo depth estimation using belief propagation. *Frontiers in Neuroscience*, 11(OCT):1–14, 2017. ISSN 1662453X. doi: 10.3389/fnins.2017.00535.

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". `classicthesis` is available for both LaTeX and LyX:

http://code.google.com/p/classicthesis/

Colors are borrowed from the *Solarized* palette, developed by Ethan Schoonover:

http://ethanschoonover.com/solarized/

Final Version as of April 16, 2018 (`classicthesis` version 1.0).