



HAL
open science

Development and validation of a Bayesian measurement technique for data-driven measurement reduction

Gonçalo Jorge Granjal Cruz

► **To cite this version:**

Gonçalo Jorge Granjal Cruz. Development and validation of a Bayesian measurement technique for data-driven measurement reduction. Fluid Dynamics [physics.flu-dyn]. Ecole Centrale de Lyon, 2024. English. NNT : 2024ECDL0012 . tel-04602281v1

HAL Id: tel-04602281

<https://theses.hal.science/tel-04602281v1>

Submitted on 12 Jun 2024 (v1), last revised 5 Jun 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



ÉCOLE
CENTRALE LYON

Numéro d'ordre : 2024ECDL0012

THÈSE DE DOCTORAT DE L'UNIVERSITÉ DE LYON

opérée au sein de

L'ÉCOLE CENTRALE DE LYON

École Doctorale 162 : Mécanique, Énergétique, Génie Civil et Acoustique

Spécialité : Mécanique des fluides et Énergétique

Soutenue publiquement le 22 Mars 2024 par :

Gonçalo Jorge GRANJAL CRUZ

Development and Validation of a Bayesian Measurement Technique for Data-Driven Measurement Reduction

Devant le jury composé de :

Christophe CORRE,
Paola CINNELLA,
James TAYLOR,
Xavier OTTAVY,
Fabrizio FONTANETO,
Juan RUIZ-SABARIEGO,
Afaf SORNIN,

Directeur de Recherche, LMFA
Professeure, Sorbonne Université
Associate Professor, Cambridge University
Directeur de Recherche CNRS, LMFA
Associate Professor, von Karman Institute
Ingénieur, Safran Aircraft Engines
Ingénieur, Safran Aircraft Engines

Présidente du Jury
Rapporteure
Rapporteur
Directeur de Thèse
Directeur de Thèse
Encadrant industriel
Examineur



ÉCOLE
CENTRALE LYON

Numéro d'ordre : 2024ECDL0012

THÈSE DE DOCTORAT DE L'UNIVERSITÉ DE LYON

opérée au sein de

L'ÉCOLE CENTRALE DE LYON

École Doctorale 162 : Mécanique, Énergétique, Génie Civil et Acoustique

Spécialité : Mécanique des fluides et Énergétique

Soutenue publiquement le 22 Mars 2024 par :

Gonçalo Jorge GRANJAL CRUZ

Development and Validation of a Bayesian Measurement Technique for Data-Driven Measurement Reduction

Devant le jury composé de :

Christophe CORRE,
Paola CINNELLA,
James TAYLOR,
Xavier OTTAVY,
Fabrizio FONTANETO,
Juan RUIZ-SABARIEGO,
Afaf SORNIN,

Directeur de Recherche, LMFA
Professeure, Sorbonne Université
Associate Professor, Cambridge University
Directeur de Recherche CNRS, LMFA
Associate Professor, von Karman Institute
Ingénieur, Safran Aircraft Engines
Ingénieur, Safran Aircraft Engines

Présidente du Jury
Rapporteure
Rapporteur
Directeur de Thèse
Directeur de Thèse
Encadrant industriel
Examineur

Ir embora pode mesmo ser a solução
Ver trabalho, brio, recompensa pela aflição
Mas se isto não mudar eu não descanso...

E se eu for
Quem te espera, Mariana?
Vais dormir nas guardas quentes de que cama?
Se eu for,
Quem vê paz na tua estampa?
Pra onde irás se eu só voltar pra pôr cá a minha campã?

Ai de mim não faço nem ideia
Prometi partir na lua cheia
Páro pra um bagaço na estação
Nos olhos de um beirão
Vejo a fera da fronteira

Diabo na Cruz. Fronteira. 2012

Acknowledgements

My first mention has to be Professor Fabrizio Fontaneto who believed in me to carry out this work from the beginning, even when I did not. My second mention goes to Professor Xavier Ottavy for supporting me on this journey, mostly from afar, but always with relevant feedback to push me forward.

I thank Safran Aircraft Engines for the opportunity given to perform this work together with the VKI and the LMFA laboratories, to which I am deeply thankful for their trust.

In these last years of my life, I had the luck to be a part of the VKI family. VKI is a great place for young researchers to learn and I hope to see it grow in the next years while keeping the same community atmosphere that I found when I first arrived in 2019 for the Research Master.

The friendships I made in that year followed me throughout my Ph.D and for that, I need to thank Giuseppe, Kayacan, Luca, Pietro, and Tommaso. It was a troubled year with COVID forcing the world to go home before the academic year finished.

When I got back to VKI to start my Ph.D. (or wait for the paperwork to go through...) I had the luck of meeting two guys doing their master's thesis who would end up joining me in the years to follow always complaining and listening to me complain about life. A special thank you to Giacomo and Lorenzo.

By the time I officially started this work, I was moved to the P116 office in the VKI Computer Center. It was there that I had the chance to meet Cedric and Riccardo on a personal level. I'm still amazed how this 3 person office, with a punk motorbike semi-pro French cyclist, a salsa-dancing fancy-dressing Italian and a Portuguese nerd like me, taught me so much.

It was around this time that I met Dr. Francisco Porres, my housemate. It is fair to say that neither of us expected to develop a brotherhood relationship like the one we did. También me presentó a las personas que se convertirían en mi familia en Bruselas: Alba, Andrea, Anita, Domenico, Elena, Joris, y Natalia. ellos hacen que la vida sea más fácil cuando llueve todos los fines de semana A todos muchas gracias.

As life continued, Natalia's role in my life grew. She helped me establish a work-life balance that I did not have while making me more self-aware about my value, directly influencing the outcomes of this manuscript. I love you. Thank you.

Si j'ai passé la majeure partie de mon temps au VKI, je n'oublierai pas les mois passés à Lyon au LMFA. C'était la que j'ai fait les experiments de validation. Merci à tous pour l'aide, specialment a Alex! Merci Christophe et Anne-Lise pour les bieres après

travail.

Esta jornada não estaria completa senão pelo apoio daqueles que não estão comigo no dia a dia mas que estão comigo quando necessário. A todas as pessoas que me acompanharam sempre na vida, mais ou menos presentes mas demasiados para nomear, um obrigado a todos, com um obrigado especial a uma engenheira brilhante que tanto fez por mim e para quem desejo tudo.

A última frase é dedicada a toda a minha família, que me apoia sempre, mesmo quando pensam que sou só doido. Nunca utilizei muitas palavras para vos explicar ou transmitir o que o vosso apoio significa para mim tão longe mas obrigado a todos do fundo do coração!

Obrigado Pai. Obrigado Mãe. Vocês são os melhores do mundo.

Résumé

Ce travail présente une méthodologie de test hybride complète pour évaluer l'écoulement dans turbomachines. Axée sur la minimisation des temps de test et des exigences en instrumentation, la méthodologie intègre de manière stratégique des mesures expérimentales standard avec des simulations numériques, en utilisant des processus gaussiens.

La méthodologie réduit systématiquement à la fois les efforts d'instrumentation et les temps de test, fournissant des métriques d'incertitude comparables aux méthodologies traditionnelles. Appliquée initialement à un compresseur axial haute pression de référence (H25) puis à un ventilateur à ultra-haut taux de dilution (ECL5 UHBR) dans des conditions de test aveugles, la méthodologie démontre sa robustesse, son adaptabilité et des réductions significatives des points de mesure et des temps de test conduisant à un impact direct sur les coûts des campagnes expérimentales.

Pour le compresseur axial H25, le cadre proposé se révèle capable de prédire les champs d'écoulement, mettant en évidence le compromis entre les mesures et l'exactitude de prédiction du flux. Les résultats du test aveugle du ventilateur ECL5 UHBR valident l'efficacité de la méthodologie dans les évaluations aérodynamiques et démontrent des économies de temps d'au moins une heure par condition de fonctionnement.

La conception d'expériences *a priori* permet une réduction d'au moins 50% des mesures, surpassant l'échantillonnage aléatoire, et assiste efficacement dans la planification de campagnes expérimentales. L'échantillonnage adaptatif *In situ* surpasse l'échantillonnage aléatoire jusqu'à 44%, démontrant une détection précise des phénomènes d'écoulement et des applications prometteuses dans la réalisation d'exigences expérimentales. La nature modulaire et adaptable de la méthodologie la positionne pour une application étendue tant dans les environnements académiques qu'industriels, tandis que son exploitation ouvre des voies pour inférer des quantités d'écoulement non mesurées ou améliorer l'évaluation des performances.

Ce travail introduit un changement de paradigme dans la planification de campagnes expérimentales, optimisant les budgets de mesure de manière stratégique à l'avance ou améliorant la précision dynamiquement au cours d'une campagne, mettant en évidence le potentiel des tendances entraînées par l'apprentissage automatique pour façonner de nouvelles voies de recherche.

Mots-clés: Technique de Mesure Hybride, Fusion de Données, Inférence Bayésienne, Apprentissage Automatique, Quantification de l'incertitude, Instrumentation, Turbo-machines

Abstract

This work presents a complete hybrid testing methodology for assessing the flow in turbomachinery components. Focused on minimizing testing times and instrumentation requirements, the methodology strategically integrates standard experimental measurements with numerical simulations, specifically employing Multi-Fidelity Gaussian Processes, Sparse Variational Gaussian Processes, and adaptive Bayesian optimization.

The methodology systematically reduces both instrumentation efforts and testing times, providing uncertainty metrics comparable to traditional methodologies. Applied initially to a benchmarked axial high-pressure compressor (H25) and afterwards to an ultra-high bypass ratio fan (ECL5 UHBR) in blind test conditions, the methodology demonstrates robustness, adaptability, and significant reductions in measurement points and testing times leading to a direct impact in experimental campaign costs.

For the H25 axial compressor, the proposed framework proves capable of predicting flow fields, emphasizing the trade-off between high-fidelity measurements and mean flow prediction accuracy. The ECL5 UHBR fan blind test results validate the methodology's efficiency in aerodynamic assessments and demonstrates time savings of at least one hour per operating condition.

The *a priori* Design of Experiments achieves at least a 50% reduction in measurements, outperforming random sampling, and effectively assists in experimental campaign planning. The *In situ* adaptive sampling outperforms random sampling by up to 44%, showcasing accurate detection of flow phenomena and promising applications in achieving high accuracy experimental demands. The modular and adaptable nature of the methodology positions it for broad application in both academic and industrial settings, while its exploitation opens paths to infer unmeasured flow quantities or improve performance evaluation measurements.

This work introduces a paradigm shift in experimental campaign planning, optimizing measurement budgets strategically beforehand or enhancing accuracy dynamically during a campaign, emphasizing the potential of machine learning-driven trends in shaping new research paths.

Keywords: Hybrid Measurement Technique, Data Fusion, Bayesian Inference, Machine Learning, Uncertainty Quantification, Instrumentation, Turbomachinery

Contents

List of Figures	xiii
List of Tables	xxi
Nomenclature	xxiii
Greek letters	xxiii
Roman letters	xxiii
Subscripts	xxiv
Operators	xxiv
Acronyms	xxiv
1. Introduction	1
1.1. Objectives of the work	5
1.2. Thesis Outline	5
2. State of the art	7
2.1. Data-Driven Approaches in Experimental Testing	8
2.2. Data Assimilation	12
2.2.1. Inverse Problems	13
2.2.2. Regression Problems	17
2.3. Bayesian Inference: Tackling Uncertainty	20
2.4. Gaussian Process Modelling	23
2.4.1. Introduction to Gaussian Process	25
2.4.2. Multi-Fidelity Gaussian Process	28
2.4.3. Stochastic Variational Gaussian Process	32
3. Hybrid Measurement Technique: Structure and Methodology	35
3.1. Introduction	36
3.2. Axial Compressor Flow Key Parameters	38
3.3. Multi-Fidelity Gaussian Process: Implementation	39
3.3.1. Kernel selection and influence	40

3.3.2.	Hyperparameters optimization	42
3.3.3.	Computational implementation	43
3.3.4.	Implementation and example: Forrester function	44
3.4.	Design of experiments: <i>A Priori</i> sampling	46
3.4.1.	SVGP as DoE tool: Implementation	47
3.5.	Adaptive <i>In situ</i> sampling	52
3.6.	Methodology Performance Evaluation	55
4.	Test Case - H25 Axial Compressor	59
4.1.	Introduction	61
4.2.	Experimental tools	62
4.3.	Numerical tools	65
4.4.	Results	67
4.4.1.	Preliminaries	70
4.4.2.	Multi-fidelity Gaussian Process	70
4.4.3.	<i>A priori</i> Design of Experiments	79
4.4.4.	<i>In situ</i> Sampling	84
5.	Test Case - ECL5 UHBR Fan	99
5.1.	Introduction	100
5.2.	Experimental tools	102
5.3.	Numerical tools	108
5.4.	Results	108
5.4.1.	Preliminaries	111
5.4.2.	<i>A priori</i> Design of Experiments	112
5.4.3.	<i>In situ</i> Sampling	126
6.	Conclusions and Future Perspectives	141
	Bibliography	147

List of Figures

1.	Introduction	1
1.1.	Evolution in engine trends over the years with crosses representing the current flying engines. Adapted from BALLAL and ZELINA, 2004. . . .	2
1.1(a).	Overall pressure ratio	2
1.1(b).	Bypass ratio	2
1.2.	Comparison of instrumentation throttling effect on high-speed and low-speed ideal fan stages.	3
1.3.	Variation of electricity prices for non-household consumers in the European Union and France from 2007 to 2023 excluding taxes. Source: EUROSTAT, 2023.	4
2.	State of the art	7
2.1.	Sketch of the reconstruction method of the non-uniform circumferential pressure signal proposed by LOU and KEY, 2021.	9
2.2.	Convergence of the spatial reconstruction temperature average with number of rakes used for different randomised rakes arrangements from SE-SHADRI, A. B. DUNCAN, et al., 2022.	11
2.3.	Schematic of weather forecast through DA adapted from ECMWF, 2017. The previous model in blue is compared to sparse observations over a set time window. The model initial state is modified through the solution of an inverse problem. The model in red shows the more accurate forecast based on DA.	14
2.4.	Variation of the CFD velocity field prediction error when increasing the number of experimental sensors used in the DA method from SOUSA, GARCÍA-SÁNCHEZ, et al., 2018.	16

2.5.	Adapted from CRUZ et al., 2021. Pressure ratio across different measurement planes. DA method matches the experimental observation at the assimilation plane. DA accurately estimates the validation experimental measurement at a different plane where the DA process was not directly influenced by the data.	17
2.6.	Normalized Google search trends evolution over time for the keywords 'machine learning' and 'data driven'. Both keywords follow a similar trend.	18
2.7.	Measured temperature field data from tomographic Schlieren integrated with a PINN to exploit physics, inferring velocity and pressure fields. Adapted from CAI et al., 2021; KARNIADAKIS et al., 2021.	20
2.8.	Schematic example of Bayesian inference. Bayesian inference refines the initial belief (prior distribution in blue) about a variable (θ) when new observed data (\mathcal{D}) is considered. The updated belief (posterior distribution in red) provides a more accurate estimate of the variable.	22
2.9.	2-dimensional multivariate Gaussian distribution. The joint probability distribution $p(\mathbf{x}_1, \mathbf{x}_2)$ is denoted at the center. The marginal distributions $p(\mathbf{x}_1)$ and $p(\mathbf{x}_2)$ are shown as projections.	25
2.10.	Classification of multi-fidelity GP approaches based on symmetric/asymmetric data fusion scheme and linear/non-linear relationship between different fidelity levels. Adapted from BREVAULT et al., 2020.	29
3.	Hybrid Measurement Technique: Structure and Methodology	35
3.1.	Proposed Data Driven Bayesian Hybrid Measurement Technique schematic for a complete accurate flow assessment and uncertainty quantification with a reduction in instrumentation usage and testing time.	36
3.2.	T-s diagram of an axial compressor stage.	38
3.3.	Prior samples for Matérn 5/2 kernel with different hyperparameter values, showcasing the effect of this parameters in GP modelling.	41
3.4.	High- and low-fidelity Forrester functions and MFGP mean prediction and uncertainty of high-fidelity function. The dots represent the available observations. In this example, 20 low fidelity and 6 high fidelity samples are used in the MFGP model which predicts accurately the behaviour of the underlying truth.	45
3.5.	Detailed schematic of the proposed <i>a priori</i> DoE flowchart with SVGP modelling and CFD simulations to determine the optimal experimental measurement locations	48

3.6.	Histogram of ELBO evaluations using batch approach. The histogram mean is presented by a red line and the true ELBO value, computed with the full gradient is presented in black.	49
3.7.	ELBO optimization process for hyperparameter and "inducing locations" selection for the Forrester function. The ELBO evolution is shown for increasing number of "inducing locations". An ELBO upper bound cluster is obtained from $N_z = 100$	50
3.8.	Detailed schematic of the proposed adaptive <i>in situ</i> DoE with the MEPE sampling approach selecting iteratively where to acquire the next experimental measurement.	56
4.	Test Case - H25 Axial Compressor	59
4.1.	H25 compressor blades	61
4.2.	VKI R4 facility layout for axial compressor studies.	63
4.3.	H25 compressor stage test section meridional view	64
4.4.	H25 experimental test grid at MP4.	65
4.5.	Miniaturized total temperature probe.	65
4.5(a).	Probe zoom.	65
4.5(b).	Probe installed in compressor stage. Blades blurred.	65
4.6.	H25 compressor stage numerical CFD domain.	66
4.7.	H25 stage outlet Total Pressure map comparison between experimental measurements (left) and CFD simulation (right).	68
4.8.	H25 stage outlet normalized Total Temperature map comparison between experimental measurements (left) and CFD simulation (right).	69
4.9.	MFGP total pressure flow assessment at MP4 (right) with different experimental measurements undersampling percentages against fully sampled experimental reference (left).	72
4.9(a).	10% Measurements used. $\rho = 0.374$	72
4.9(b).	20% Measurements used. $\rho = 0.158$	72
4.9(c).	33% Measurements used. $\rho = 0.110$	72
4.10.	RMSE of Total pressure ratio of across 50 MFGP models. Sensitivity to the random selection of experimental data and the percentage of experimental data used in the model training.	74
4.11.	Total pressure ratio uncertainty prediction for MFGP model presented in Figure 4.9c with 33% measurements used.	77
4.11(a).	Total pressure flow.	77

4.11(b).	Uncertainty.	77
4.12.	MFGP total temperature flow assessment at MP4 (right) with 33% experimental measurements undersampling against fully sampled experimental reference (left) with $\rho = 0$	78
4.13.	ELBO optimization for hyperparameter and 'inducing locations' selection for the H25 test case CFD RANS simulation. A cluster of the ELBO from around $N_z = 300$ data points sets the DoE number of measurements required.	81
4.14.	Nearest neighbour effect on the DoE optimal sampling grid. SVGP selected locations and the NN approximate location used, in green and red respectively.	82
4.15.	RMSE comparison between random sampling, the <i>a priori</i> SVGP DoE proposed sampling approach with NN and without NN across 50 model runs for different amount of measurements used.	83
4.16.	Proposed adaptive sampling performance evaluation and error/uncertainty weight criteria for pressure ratio flow field assessment with increasing number of measurements.	86
4.16(a).	RMSE	86
4.16(b).	Maximum Uncertainty	86
4.16(c).	Weight criteria α	86
4.17(a).	5% measurements used. RMSE $\approx 5.8\text{E-}03$	88
4.17(b).	15% measurements used. RMSE $\approx 1.2\text{E-}03$	89
4.17(c).	20% measurements used. RMSE $\approx 8.8\text{E-}04$	89
4.17(d).	33% measurements used. RMSE $\approx 3.9\text{E-}04$	90
4.18.	RMSE comparison between random, <i>a priori</i> SVGP DoE and adaptive sampling approaches across 50 model runs for different amount of measurements used.	92
4.19.	RMSE comparison between different initial datasets for the proposed adaptive sampling across 50 model runs for different amount of measurements used.	94

4.20.	Pressure flow field assessment with the proposed adaptive sampling DoE with the defined stopping criteria met. On the top left, the reference experimental flow and the measurement samples used are shown. On the top right, the GP model mean flow prediction is presented. The bottom left shows the absolute difference between the experimental reference and the predicted pressure ratio flow. The bottom right contour displays the predicted GP model flow uncertainty.	96
5.	Test Case - ECL5 UHBR Fan	99
5.1.	ECL5 rotor blade and fan assembly adapted from SCHNEIDER et al., 2023.	101
5.1(a).	ECL5 blade.	101
5.1(b).	Fan assembly.	101
5.2.	ECL-B3 rig schematic.	102
5.3.	ECL5 fan stage meridional view schematic.	103
5.4.	L-shape bent five hole probes (5HP) deployed at MP 'SE' for flow assessment based on the proposed methodology.	104
5.4(a).	Hub probe zoom.	104
5.4(b).	Hub and casing 5HP.	104
5.4(c).	Casing probe zoom.	104
5.5.	End-wall proximity of the 5HP for both hub and casing with measurements possible up to 1.7% and 98.3% of the span, respectively.	105
5.5(a).	Hub.	105
5.5(b).	Casing.	105
5.6.	ECL5 experimental test grid with black and blue dots representing casing and hub probe respectively.	105
5.7.	Reference test pressure field contour obtained with the two L-shaped 5HP in parallel with an overlap region.	106
5.8.	ECL5 Total Pressure map comparison between experimental measurements (left) and CFD simulation (right) for the OP at nominal rotational speed.	110
5.8(a).	100% rotational speed PP.	110
5.8(b).	100% rotational speed DP.	110
5.8(c).	100% rotational speed NC.	110

5.9.	ELBO optimization for hyperparameter and 'inducing locations' selection for the ECL5 test case CFD RANS simulation for DP at 100% rotational speed. A cluster of the ELBO from around $N_z = 400$ data points sets the DoE number of measurements required.	114
5.10.	ECL5 mean total pressure and uncertainty contours for DP at 100NM inferred with <i>a priori</i> SVGP DoE and MFGP model. Comparison to experimental reference.	116
5.10(a).	Mean flow assessment (right) against fully experimental reference (left).	116
5.10(b).	Uncertainty.	116
5.11.	Absolute pressure difference between reference experimental data and the inferred flow field, highlighting local error regions.	118
5.12.	ECL5 mean total pressure and uncertainty contours for PP at 100NM inferred with <i>a priori</i> SVGP DoE and MFGP model. Comparison to experimental reference.	119
5.12(a).	Mean flow assessment (right) against fully experimental reference (left).	119
5.12(b).	Uncertainty.	119
5.13.	ECL5 mean total pressure and uncertainty contours for NC at 100NM inferred with <i>a priori</i> SVGP DoE and MFGP model. Comparison to experimental reference.	120
5.13(a).	Mean flow assessment (right) against fully experimental reference (left).	120
5.13(b).	Uncertainty.	120
5.14.	Pressure ratio pitch comparison of the different kernel at for a fixed $r = 0.63$	124
5.15.	Initial set of $N_{y_{m0}} = 40$ measurements obtained with optimal LHS for the ECL5 test case.	127
5.16.	Final $N_y = 220$ measurements obtained with the adaptive sampling blind test for the ECL5 test case.	128
5.17.	ECL5 lower radius pressure field and uncertainty assessed with the GP model with the adaptive sampling measurements of the hub probe. . .	129
5.17(a).	Mean flow assessment.	129
5.17(b).	Uncertainty evaluation.	129
5.18.	ECL5 pressure field and uncertainty evaluated with the MFGP model after the adaptive sampling measurements of the hub probe.	131
5.18(a).	Mean flow assessment.	131
5.18(b).	Uncertainty evaluation.	131

5.19.	Simulated ECL5 adaptive sampling error/uncertainty and weight criteria for pressure ratio flow field assessment with increasing number of measurements.	133
5.19(a).	RMSE	133
5.19(b).	Maximum Uncertainty.	133
5.19(c).	Weight criteria α	133
5.20(a).	120 measurements acquired. RMSE $\approx 2.2\text{E-}02$	135
5.20(b).	220 measurements acquired. RMSE $\approx 1.3\text{E-}02$	136
5.20(c).	240 measurements acquired. RMSE $\approx 8.3\text{E-}03$	136
5.20(d).	360 measurements acquired. RMSE $\approx 6.0\text{E-}03$	137
6.	Conclusions and Future Perspectives	141
6.1.	MFGP tangential pressure profile inference at the H25 compressor probe sampled limits.	143
6.2.	MFGP model compressor map using mass-averaged CFD simulations and mass-averaged rake measurements.	144
	Bibliography	147

List of Tables

4. Test Case - H25 Axial Compressor	59
4.1. H25 stage design parameters.	61
4.2. Experimental hyperparameters and performance metrics for the different models presented in Figure 4.9.	73
4.3. RMSE sensitivity analysis for the H25 test case across 50 model runs different amount of measurements used.	75
4.4. MaxAE sensitivity analysis for the H25 test case across 50 model runs with different amount of measurements used.	75
4.5. RMSE sensitivity analysis with proposed <i>A priori</i> SVGP DoE across 50 different runs for a different amount of measurements used.	84
4.6. RMSE sensitivity analysis with proposed <i>A priori</i> SVGP DoE across 50 different runs for a different amount of measurements used.	92
5. Test Case - ECL5 UHBR Fan	99
5.1. ECL5 stage design parameters.	100
5.2. ECL5 experimental operating points (OP).	107
5.3. Number of optimal 'inducing locations' N_z for all OP tested in the ECL5 fan test case.	114
5.4. Quantitative evaluation of the relevant estimated hyperparameters and the error metrics for the MFGP models for the different ECL5 OPs.	118
5.5. Comparison of error metrics between the different kernels for the pressure field inference of the ECL5 fan at DP for 100% rpm.	123

Nomenclature

Greek letters

δ	MFGP bias/variation
η	Observation noise
γ	Specific heat ratio
ι	Isentropic Efficiency
μ	Mean quantity
Π	Total to Total Pressure Ratio
ρ	Scaling factor
Σ	Covariance
σ	Variance
θ	Parameter

Roman letters

\mathcal{D}	Dataset
f	Generic function
\mathcal{G}	Forward model
\mathbf{I}	Identity Matrix
\mathbf{K}	Kernel
k	Data dimensions
\mathbf{L}	Matrix transformation
ℓ	Length scale
\dot{m}	Mass Flow Rate [kg/s]
N	Rotational Speed
n	Data size

p	Probability
q	Variational distribution
s	Entropy [J/K]
T	Temperature [K]
X	Input variable
y	Observed quantity
Z	Inducing points

Subscripts

\cdot_*	Predicted quantity
\cdot_f	Kernel scaling factor variance
\cdot_{high}	High fidelity
\cdot_{HS}	High speed
\cdot_{in}	Inlet
$\cdot_{i,j}$	Indexes
\cdot_{low}	Low fidelity
\cdot_m	Adaptive sampling iteration
\cdot_{max}	Maximum quantity
\cdot_{min}	Minimum quantity
\cdot_{out}	Outlet
\cdot_t	Fidelity level
\cdot_0	Total quantity
\cdot_y	GP model variance
\cdot_z	Inducing locations

Operators

\cdot^{-1}	Inverse operator
$ \cdot $	Euclidean norm operator in \mathbb{C}
\cdot^T	Transpose operator

Acronyms

5HP 5-Hole Probe

AR1	Auto-Regressive Linear Method
BO	Bayesian Optimization
CFD	Computational Fluid Dynamics
DA	Data Assimilation
DNS	Direct Numerical Simulations
DoE	Design of Experiments
DP	Design Point
ECMWF	European Centre for Medium-Range Weather Forecasts
ELBO	Evidence Lower Bound
EPE	Expected Predicted Error
GP	Gaussian Process
IEnKF	Inverse Ensemble Kalman filter
KL	Kullback-Leibler divergenc
L-BFGS-B	Limited-memory Broyden-Fletcher-Goldfarb-Shanno with Bounds
LHS	Latin Hypercube Sampling
LMFA	Laboratoire de Mécanique des Fluides et d'Acoustique
LOOCV	Leave-One-Out Cross-Validation
MEPE	Maximize Expected Predicted Error
MF-DGP	Multi-Fidelity Deep Gaussian Process
MFGP	Multi-Fidelity Gaussian Process
ML	Machine Learning
NARGP	Non-linear Auto-Regressive multi-fidelity Gaussian Process
NC	Near-Choke Pressure Ratio Point
NN	Nearest Neighbour
OP	Operating Point
e1sA	<i>Ensemble Logiciel de Simulation Aérodynamique</i>
PINN	Physics-Informed Neural Network
PIV	Particle Image Velocimetry
POD	Proper Orthogonal Decomposition
PP	Maximum Pressure Ratio Point
RANS	Reynolds-Averaged Navier-Stokes

RBF	Radial Basis Functions
RMS	Root Mean Square
SAE	Safran Aircraft Engines
SGD	Stochastic Gradient Descent
SVGP	Sparse Variational Gaussian Process
UQ	Uncertainty Quantification

Chapter 1.

Introduction

AIRCRAFT engine manufacturers are forced to deal with challenges arising from the intersection of stringent government regulations and environmental concerns. These challenges manifest in the form of emissions regulations and constraints on fuel consumption.

In response to these challenges, the pursuit of enhanced engine performance has emerged as a pivotal objective, defining industry leaders. This pursuit centres on the critical goal of reducing fuel consumption, emissions, and noise. One strategy involves the enhancement of the Brayton cycle efficiency by increasing the overall pressure ratio at the compressor outlet, for which, a solution, is to reduce the channel core dimensions and thus, control the tip Mach number. This reduction leads to a subsequent decrease in compressor blade height. Another approach to tackle these issues involves increasing the engine bypass ratio through the introduction of a large-diameter geared low-speed fan. This enhances propulsive efficiency by reducing outlet jet velocity, resulting in reduced noise emissions and decreased tip speeds. Figures 1.1a and 1.1b demonstrate, respectively, the evolution of engine pressure and bypass ratios over time, showcasing this trend. Notably, current flying engines from major manufacturers are indicated with crosses in the figures, illustrating the ongoing efforts in this direction.

In the frame of the development of the next-generation engine concepts, that continue the highlighted trend, a need emerges for a shift in conventional experimental testing methodologies. Recent years have underscored the necessity for an evolution in testing practices to align with recent literature trends. This evolution calls for new experimental methodologies to overcome standard measurements techniques limitations.

High pressure compressors have seen a drastic reduction in blade height. Instrumentation intrusiveness induces an increased blockage and an inevitable variation of the local flow field with a direct impact on the flow assessment. The instrumentation intrusiveness becomes apparent when evaluating these machines performance.

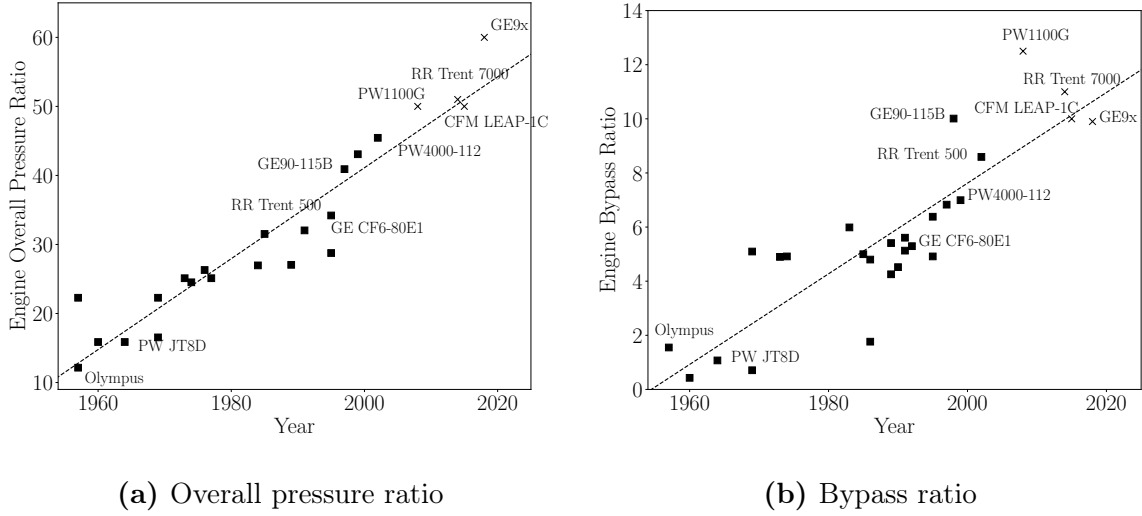


Figure 1.1. Evolution in engine trends over the years with crosses representing the current flying engines. Adapted from BALLAL and ZELINA, 2004.

The heightened sensitivity to instrumentation extends beyond small engine components to geared low-speed fans. A reduction in the pressure ratio generated by a fan stage inherently affects the total temperature ratio. The assumption of constant instrument accuracy results in a noticeable decline in performance measurement precision, particularly concerning stage efficiency.

Figure 1.2 compares the ideal loading characteristic curves of two fan stages operating at different rotational speeds, high speed (HS) and low speed (LS) normalized with respect to design conditions. Considering a standard high-speed fan, for a fixed instrumentation blockage, a throttling effect $\delta\phi$ on the operating condition is given by δ_{HS} . If in a state of the art low-speed fan the throttling effect $\delta\phi$ is kept constant, the loading variation δ_{LS} is clearly higher than in its predecessor ($\delta_{LS} > \delta_{HS}$), thus imposing a larger variation of the operating point just from a instrumentation intrusiveness point of view. Essentially, instrumentation becomes a direct source of experimental error, significantly impacting performance characterization, especially in the characteristic maps near the stall region.

In response to the heightened sensitivity of instrumentation, researchers have sought innovative solutions that minimize interference. In recent years, non-intrusive optical measurement techniques, such as Particle Image Velocimetry (PIV) and Laser Doppler Anemometry (LDA) have gained attention. These techniques offer the advantage of capturing data without perturbing the flow and have had considerable success in controlled environments and lab-scale experiments.

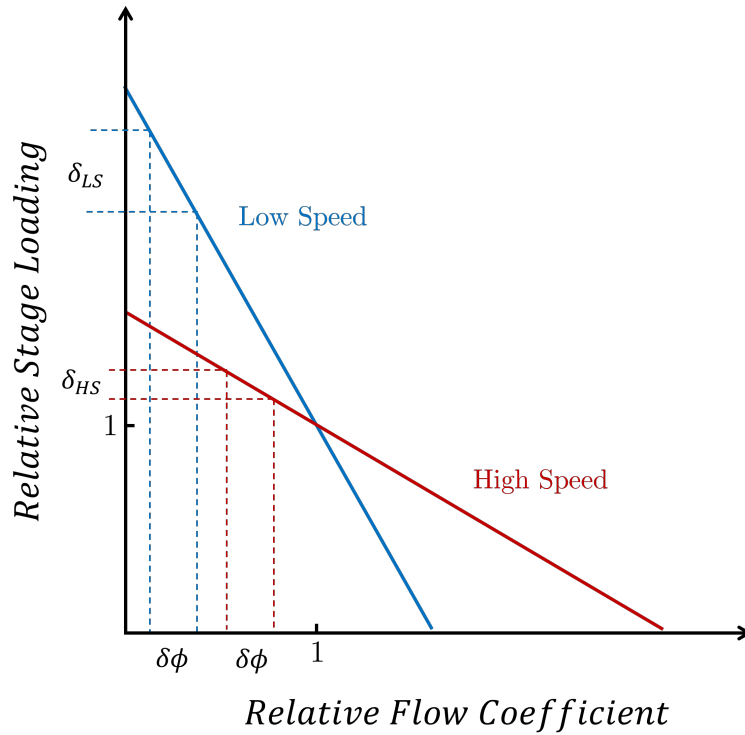


Figure 1.2. Comparison of instrumentation throttling effect on high-speed and low-speed ideal fan stages.

Simultaneously, a dedicated effort aims to push the boundaries of instrumentation miniaturization. Researchers are developing smaller probes to minimize interference with flow dynamics. However, this faces limitations. The reduction of probe size comes with inherent trade-offs, including diminished bandwidth resolution and challenges in maintaining structural integrity.

While these approaches are areas of active research, translating them to real engine testing environments presents significant hurdles. Non-intrusive techniques require flow optical access and not all testing facilities possess the infrastructure to accommodate these needs. Moreover, in a full-scale engine environment these techniques have yet to be applied due to the harsh temperature and pressure, but also contaminants that hinder the optical systems' durability and accuracy, thus making them an impractical solution within an industrial context.

The ongoing miniaturization of instrumentation probes approaches an asymptote. Probes cannot be infinitely reduced in size without compromising their functionality and accuracy. This constraint is particularly pronounced in harsh operating conditions where both precision and robustness are essential. Additionally, specially in an industrial context where full engine components are tested, there is an increased difficulty

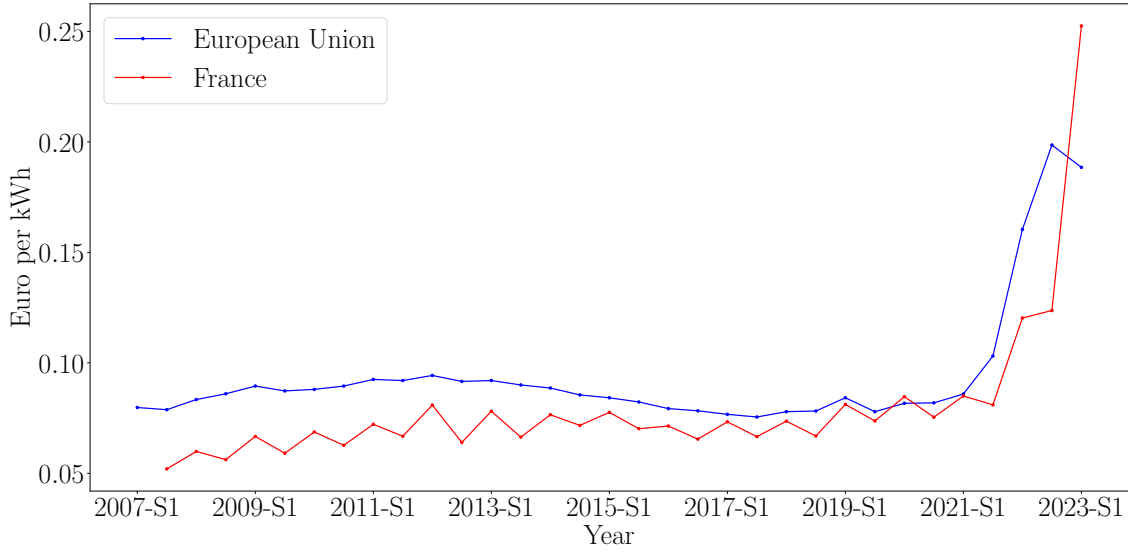


Figure 1.3. Variation of electricity prices for non-household consumers in the European Union and France from 2007 to 2023 excluding taxes. Source: EUROSTAT, 2023.

in instrumenting specific regions of the flow inside the machine mainly due to access constrains.

Amidst these challenges, an additional major factor commands the industry attention in the present day when planning experimental measurement campaigns: the escalation of energy prices.

The recent increase in energy costs, showcased in Figure 1.3 for the European Union and France electricity price from 2007 to 2023, adds complexity to the planning of an experimental measurement campaign with the electricity price increasing more than two-fold in the last two years. This increase significantly contributes to higher operational costs during testing, making testing time a critical budget factor.

With the decreasing size of the new generation compressor core or the design of new low-speed geared ultra-high bypass ratio fans, the demand for comprehensive flow sampling across various operating points and rotational speeds intensifies, aiming to gain a profound understanding of flow phenomena. Testing time thus becomes a critical budget factor, especially in an industrial context where full engine components are tested in Megawatt powered rigs. Prolonged testing periods substantially contribute to higher operational costs, driven by escalating energy expenses.

Both researchers and industry are compelled to decrease the test matrix, the sampled operating points and the spatial acquisition discretization. While this addresses energy-

related costs to some extent, it compromises the comprehensiveness of flow assessment, potentially yielding incomplete insights.

1.1. Objectives of the work

Building upon these challenges, the main objective of this work is to propose a novel hybrid testing methodology for the performance characterisation and flow assessment of turbomachinery components. This methodology aims to address the limitations posed by instrumentation while significantly reducing testing times. The proposed methodology objectives are the following:

1. The methodology should employ a hybrid data-driven approach, which combines experimental techniques with numerical models.
2. The experimental techniques should leverage standard instrumentation technology and the methodology must have a low computational cost, ensuring the methodology's widespread practicality.
3. The methodology should include a propagated uncertainty quantification linked to its specific application and the assessed flow quantities.
4. The methodology should possess predictive capabilities, enabling its robust application within a design of experiments (DoE) framework that enables the design of an experimental campaign beforehand significantly influencing the efficiency and cost of the measurement campaign.
5. The methodology must be general, independent of specific test cases, facilities, or flows.

By redefining how experimental campaigns are conducted, this study aims to support advancements in aviation propulsion testing.

1.2. Thesis Outline

The contents of this thesis are divided in six chapters, including this introduction, providing the motivation for the work, followed by its objectives.

The second chapter presents the literature review, which aims to provide the theoretical ground for methodology proposed.

The third chapter introduces the methodology, presenting its workflows and discussing crucial implementation choices. Their influence on the results are also discussed.

The fourth chapter presents the results obtained for the proposed methodology applied to an axial high-pressure compressor. Starting with a description of the experimental and numerical tools used, the presentation and analysis of the results. In fifth chapter, the proposed methodology is applied to a low-speed ultra high bypass ratio fan.

Lastly, the sixth chapter concludes this work, highlighting the implications of this work for the field. An outlook discussion of the future work strategies to further improve the methodology follows.

Chapter 2.

State of the art

THIS chapter outlines the available literature and provides the necessary mathematical foundation of an experimental data-driven hybrid approach for flow and performance assessment. The hybrid methodology introduced later in this work is contextualized by examining the literature of data assimilation, uncertainty quantification, and Gaussian process modelling.

Overview

2.1. Data-Driven Approaches in Experimental Testing	8
2.2. Data Assimilation	12
2.2.1. Inverse Problems	13
2.2.2. Regression Problems	17
2.3. Bayesian Inference: Tackling Uncertainty	20
2.4. Gaussian Process Modelling	23
2.4.1. Introduction to Gaussian Process	25
2.4.2. Multi-Fidelity Gaussian Process	28
2.4.3. Stochastic Variational Gaussian Process	32

2.1. Data-Driven Approaches in Experimental Testing

The increasing complexity of next generation turbomachinery components sets the context for a need of accurate flow and performance assessment. Traditional experimental methodologies often involve time-consuming processes and intrusive instrumentation methods that alter the flow characteristics.

Optical techniques offer advantages such as non-intrusiveness, high resolution, and the ability to capture transient phenomena. For instance, YU and B. J. LIU, 2007 employed stereo-PIV to study the rotor tip-flow secondary structures in a low speed compressor stage. ZAMBONINI et al., 2017 utilized PIV to analyse the 3D unsteady behaviour of a compressor corner separation, although in a linear cascade configuration. More recently, OKADA et al., 2024, conducted PIV measurements in a low-pressure turbine cascade while testing at relevant engine conditions. However, the applicability of these techniques within an industrial context, especially at higher Technology Readiness Levels (TRL), remains an open question.

Several data-driven experimental approaches have emerged to deal with intrusiveness of the instrumentation while guaranteeing an accurate flow assessments. These approaches leverage extensive datasets to reconstruct missing or inaccessible data such as done by VENTURI and KARNIADAKIS, 2004, where Proper Orthogonal Decomposition (POD) is used to reconstruct 'gappy' direct numerical simulations (DNS) snapshots in an accurate spatio-temporal velocity fields. Their robustness analysis, assessing the required amount of data for an accurate reconstruction, revealed that noise in the flow assessment arises when dealing with undersampled data at around 80% of the complete dataset. Despite the initial application to DNS data, this work was a precursor and POD has been accurately applied to improve the reconstruction of PIV data by RAIOLA et al., 2015, showing the applicability of data driven techniques in experimental fluid mechanics.

While POD data-driven approach has its value, it also presents its challenges. A major drawback is the need of large amounts of data, that are hard to obtain in a turbomachinery industrial facility due to high experimental costs and difficult flow access. To overcome the need for large amounts of data, some authors have relied on fusing different data sources, for example, BERTRAM et al., 2021 extend the usage of data heavy POD to fuse different data sources and support pressure taps experiments with an extensive numerical database over a flight envelope to study the pressure distribution in an aircraft wing.

In the field of turbomachinery, one example is the work of LOU and KEY, 2021;

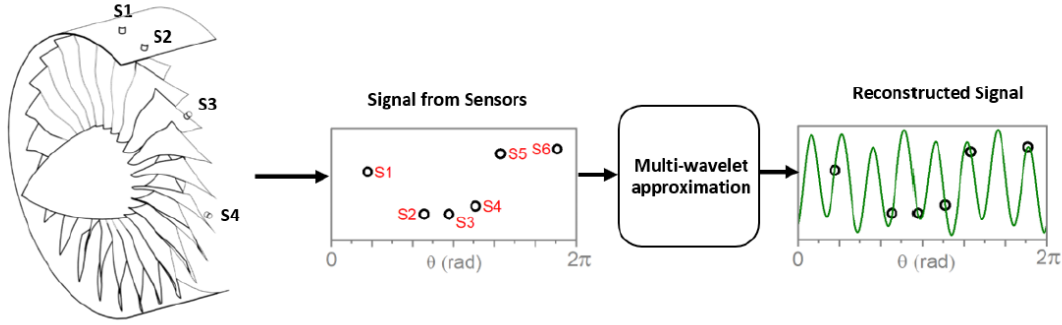


Figure 2.1. Sketch of the reconstruction method of the non-uniform circumferential pressure signal proposed by LOU and KEY, 2021.

LOU, MATTHEWS, et al., 2021, where a non-uniform circumferential pressure signal is reconstructed from spatially undersampled data in a multistage axial compressor. This reconstruction is sketched in Figure 2.1. Acquired pressure signals at different tangential locations along the annulus are employed in multi-wavelet approximation method to reconstruct a full annulus pressure signal. To accurately capture the compressor’s critical wavenumbers, an optimization algorithm determines the optimal placement of pressure probes. The optimization function is subject to geometric constraints, which are common in experimental test facilities using casing-mounted probes. The authors included a minimum spacing between adjacent probes, defined as:

$$|\Delta\theta_{j,i}| = |\theta_j - \theta_i| \geq \theta_{\min} \quad (2.1)$$

where here θ is the probe’s circumferential position. Furthermore, obstructions on the facility’s casing/test section fixtures are addressed by the following constraint:

$$\theta_i \in [\theta_{1,\min}^*, \theta_{1,\max}^*] \cdots [\theta_{m,\min}^*, \theta_{m,\max}^*] \quad (2.2)$$

where each m range corresponds to regions with obstructed access in the facility.

The authors in LOU and KEY, 2021 emphasize the broader implications of their analysis. They underline that their approach not only addresses the challenge of reconstructing circumferentially non-uniform flow but also contributes to improving the averaging and performance evaluation of the machine against standard rake area-averaging procedures that might fail to quantify circumferential non-uniformities. CHILLA et al., 2020 initially carried out an investigation on instrumentation errors attributed to circumferential flow variations in the same multi-stage axial compressor and demonstrated

that a baseline probe configuration with three equally spaced probes around the annulus could result in a maximum error of 2.8 points error in compressor efficiency.

Traditionally, experimental characterization of circumferential flow variations is achieved through circumferential traversing of probes on stator rows, necessitating complex traverse mechanisms that are cumbersome to operate and incur high costs associated with extensive testing times.

Another significant contribution that addresses the challenge of accurate averaging while assessing flow characteristics is found in the research by SESHADRI, A. DUNCAN, et al., 2022. In their study, the authors propose a statistical methodology designed to refine the process of flow averaging, specifically when working with limited rake temperature measurements. They claim that the proposed averaging is more representative than traditional rake area averaging and a reduced number of rakes is required.

This innovative methodology builds upon a data-driven Bayesian modelling previously developed by the authors in SESHADRI, A. B. DUNCAN, et al., 2022, in which, the authors approximate the spatial temperature flow field 'primary aerothermal features'. This approximation is achieved using a limited number of thermocouple measurements placed within rakes along the annulus within a Bayesian framework that is explained and exploited later in this work. The modelling process involves employing a Fourier basis to effectively capture the circumferential variation with fixed harmonics. To validate their approach, the authors conduct a meticulous comparison between their spatial reconstruction and an engine test. In the engine test, rakes are tangentially traversed along the full annulus, enabling a comprehensive sensitivity analysis encompassing different rake arrangements.

Figure 2.2 visually captures the outcome of this sensitivity analysis. It showcases the convergence behaviour of the average temperature as the number of rakes employed increases. It depicts various random rake arrangements for each number of rakes, illustrating the impact of the rake arrangement positions on the correct evaluation of the temperature. It becomes evident that while an increased number of rakes naturally leads to better assessment, an optimal arrangement of even a small number of rakes would significantly influence the quality of the evaluation achieved.

These conclusions are pertinent to the current work and underscore the critical role that experimental setup and instrumentation play in accurately assessing the flow field's behaviour. The amount of data available and the location where it is obtained are paramount for the correct application of data-driven techniques in experiments.

While data-driven experimental approaches have proved their value in accurately assessing flow behaviour and machine performance, they simultaneously bring forth the

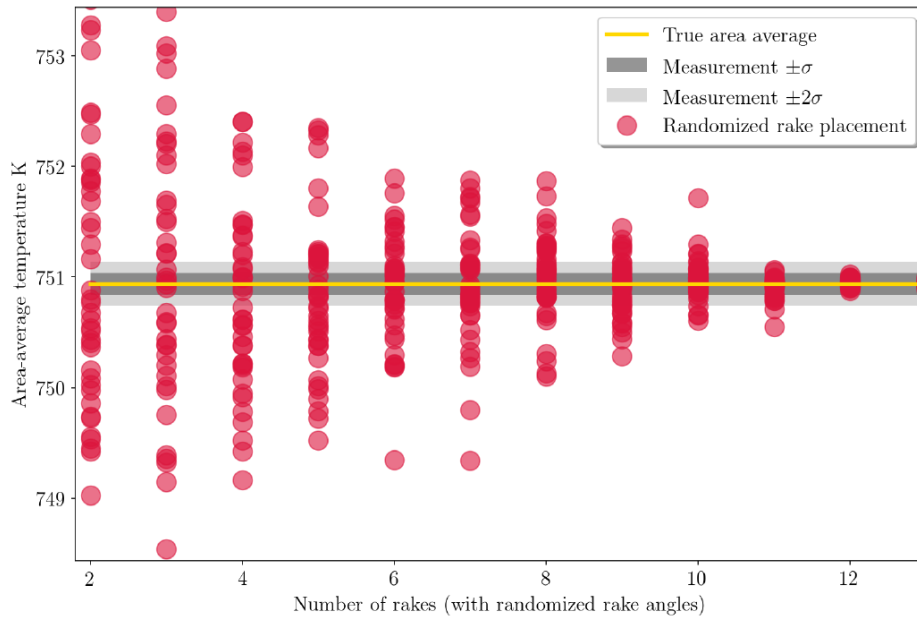


Figure 2.2. Convergence of the spatial reconstruction temperature average with number of rakes used for different randomised rakes arrangements from SEHADRI, A. B. DUNCAN, et al., 2022.

challenges of data availability and the balance between too much and little data.

On one hand, techniques like those proposed by VENTURI and KARNIADAKIS, 2004 demand substantial amounts of data, a requirement often met with difficulties and high costs, especially in turbomachinery industrial settings. On the other hand, relying solely on sparse yet strategically placed measurements, as demonstrated by LOU and KEY, 2021 and SESHADRI, A. DUNCAN, et al., 2022, can risk compromised accuracy and result in uncertainty propagation. The balance between data abundance and data optimization prompts the exploration of innovative approaches that leverage both comprehensive and sparse data. This is where the concept of data assimilation comes into play.

In the following section, data assimilation is introduced, contextualizing how this approach can be employed in an hybrid measurement technique that bridges between different data sources to improve the precision of flow and performance evaluation in turbomachinery components while reducing the instrumentation effort and extensive testing times.

2.2. Data Assimilation

In the last decades, the remarkable advancements in computational power have established computational fluid dynamics (CFD) simulations as a standard tool, enabling researchers to gain insights into the complete flow field domain and unsteady phenomena.

While the prevalence of CFD methods might lead to the assumption that certain traditional experiments are now obsolete due to cost and sparse information, the reality has proven otherwise. Even as CFD simulations continue to evolve, issues such as numerical discretization and turbulence modelling remain open challenges topics in the literature. Furthermore, modern turbulence modelling, which play a critical role in CFD, relies heavily on high-precision experimental data to achieve accurate predictions, with SINGH et al., 2017 proposing a data driven approach to fine tune turbulence models, or in high-fidelity DNS data, with SCHMELZER et al., 2020 directly inferring closure equations from it.

The increased availability of CFD data enables innovative data-driven approaches that combine computational models with real-world measurements. This domain is known as data assimilation (DA), which fundamentally involves the synergy of numerical simulations and observational data, as presented in ASCH et al., 2016. In the context of fluid mechanics, DA offers a means to seamlessly integrate experimental data with CFD simulations.

While experimental measurements provide accuracy, they are constrained in terms of spatial and temporal coverage. On the other hand, CFD simulations offer a complete domain analysis, but still struggle with the challenge of accurately representing complex real-world flows due to numerical instabilities, boundary condition inaccuracies, and turbulence modelling that lead to hard to estimate epistemic errors and high uncertainties. Consequently, DA offers a solution to address two distinct challenges: the potential discrepancy between CFD predictions and real-world flow states, and the constraints imposed by limited spatial and temporal coverage in accurate experimental data HAYASE, 2015.

The first challenge to tackle the discrepancies between CFD predictions and real-flow conditions aligns with the mathematical formulation of inverse problems. Inverse problems involve inferring parameters by leveraging available observations, thus 'inverting' typical forward modelling approaches. In fluid mechanics, this usually translates into improving boundary conditions, or turbulence model parameters to guide CFD simulations into results consistent with observed experimental data. DA techniques such as

ensemble methods and variational inference are usually selected to continuously adjust parameters and improve the match between simulation and experimental observations.

Addressing the second challenge of dealing with limited spatial or temporal sampling can be mathematically formulated as a regression problem. Regression algorithms aim at uncovering patterns and establishing relations based on available data. In fluid mechanics, they are often applied to develop flow fields models that interpolate or extrapolate available finite data. This is the main challenge addressed in this work.

The following subsections explore these mathematical formulations - inverse and regression problems - within the literature context of building towards data-driven approaches for enhancing experimental fluid mechanics methods.

2.2.1. Inverse Problems

Inverse problems form a fundamental mathematical framework that establishes a link between different data sources, such as computational fluid dynamics (CFD) predictions and real-world flow states, through DA. The concept of an inverse problem is to retrieve missing information through a model whose input parameters are defined based of a limited number of observations.

Mathematically, this concept materializes in finding the appropriate parameters $\boldsymbol{\theta}$ given a vector of observed data \mathbf{y} , as expressed by the equation:

$$\mathbf{y} = \mathcal{G}(\boldsymbol{\theta}) + \eta \quad (2.3)$$

in this formulation, the function \mathcal{G} represents the forward mapping of the unknown input parameters $\boldsymbol{\theta}$ through a simulation model, and η is the noise inherent in the observation of \mathbf{y} .

The inverse problems framework has found extensive use in the field of weather forecasting. In this context, DA enables meteorologists to merge real-time observations with predictive models, allowing for more accurate weather forecasts and aiding in the understanding of complex atmospheric dynamics. This is performed with experimental data from satellite information or ground measurements that continuously are acquired and used to improve weather forecasting models. This is done, for instance, by the European Centre for Medium-Range Weather Forecasts (ECMWF) which applies DA to improve weather predictions. Figure 2.3 illustrates an example where DA is used over a specific time window to improve a model forecast of a single quantity x . Based on the available observations and model, the initial model state is modified to achieve an improved model that provides a better match to experimental observations and a




Figure 2.3. Schematic of weather forecast through DA adapted from ECMWF, 2017. The previous model in blue is compared to sparse observations over a set time window. The model initial state is modified through the solution of an inverse problem. The model in red shows the more accurate forecast based on DA.

more accurate weather prediction ECMWF, 2017.

The inverse problems literature usually divides in two approaches: variational methods and ensemble methods, or an hybrid of the two. All the methods leverage Bayesian statistics, a fundamental statistical theory that will be later described in detail. BANISTER, 2017; LAW and STUART, 2012 provide a good review into these methods and their context in the literature, while VETRA-CARVALHO et al., 2018 give a coherent mathematical description of the different methods and provide important implementation guidelines and approaches to enforce regularization of ill-posed problems.

Variational DA tackles the inverse problem as an optimization problem, aiming to minimize a cost function. It constantly updates the model by seeking to minimize the difference between observed data and model predictions by iteratively adjusting input parameters.

One example of variational DA to fundamental fluid mechanics can be found in the work of FOURES et al., 2014. They use variational DA to assimilate Reynolds-averaged Navier-Stokes (RANS), with the Reynolds stresses chosen as the assimilation parameter. Their algorithm was able to find the optimal forcing and reconstruct the full mean flow field around a cylinder at a Reynolds number of 150 to match DNS data.

SYMON et al., 2017 later proved the approach viable for a high Reynolds number with time-averaged PIV experimental data over an airfoil, thus obtaining a mean-flow that is not limited by the PIV field of view and low spatial discretization. To achieve it, the original framework was modified, with a smoothing process being added to it for regularization in all domain points. Regularization involves the introduction of additional information or constraints into the problem formulation, which aids in stabilizing the solution process. This additional information can come in various forms, such as prior knowledge about the parameters or assumptions about the smoothness of the solution.

Ensemble DA methods are based on the ensemble Kalman filter made popular by EVENSEN, 2010. Its essence lies in the creation of an ensemble of potential parameter sets, spanning the parameter space. Each ensemble member serves as a representative model. These methods function iteratively, driven by the information available from measurements. The refined parameters enable statistical computations that take into account the uncertainties present in the available data. This ensemble-based approach not only provides insight into the unknown parameters but also offers a probabilistic distribution perspective.

In the context of urban flows, SOUSA, GARCÍA-SÁNCHEZ, et al., 2018; SOUSA and GORLÉ, 2019 employ a DA ensemble method, the inverse ensemble Kalman filter (IEnKF) IGLESIAS et al., 2013, to address the complexity and variability of the inlet atmospheric boundary layer in the predictive accuracy of CFD models. The IEnKF is used to assimilate data from optimally deployed anemometers within the campus of the university of Stanford campus to correctly estimate the inflow boundary conditions for RANS simulations. By integrating the IEnKF, they significantly enhance the accuracy of RANS predictions, surpassing traditional methods that rely on weather station data for boundary condition definition.

The significance of this application to the present work is the fact that the accuracy and reliability of the predictions are influenced by the number and location of sensors used for assimilation. Figure 2.4 illustrates the impact of adding more assimilation sensors to an initial single sensor with artificially added high noise on CFD prediction error. As the number of sensors grows, the overall prediction error decreases, effectively compensating for the single sensor high noise, specially in flow regions where the sensors are deployed. This finding aligns with the aim of streamlining instrumentation and reducing the number of measurements while ensuring robust predictions.

A parallel example of ensemble DA's potential is evident in the field of turbomachinery, as demonstrated by CRUZ et al., 2021. This work focuses on inferring boundary

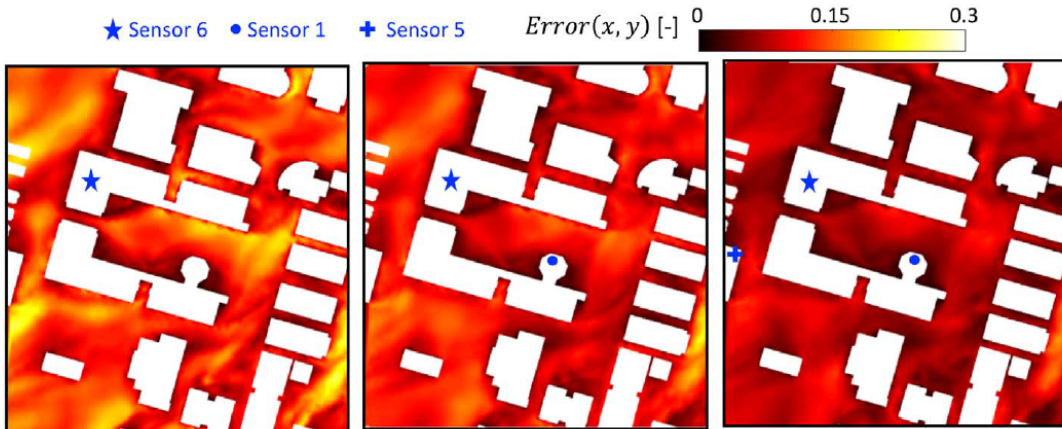


Figure 2.4. Variation of the CFD velocity field prediction error when increasing the number of experimental sensors used in the DA method from SOUSA, GARCÍA-SÁNCHEZ, et al., 2018.

conditions for CFD RANS simulations of an axial compressor. This includes operating parameters such as inlet total pressure, inlet total temperature, and mass flow. The flexibility of the ensemble DA algorithm allows to extract additional information, such as inter-row pressure, which was not part of the assimilation data.

Figure 2.5 adapted from CRUZ et al., 2021 showcases their approach. The diagram illustrates the axial compressor’s measurement planes. Plane 0 represents the compressor inlet, where initial boundary conditions are set. Plane 2 corresponds to the compressor inter-row, where no experimental data were assimilated. Plane 4 is the location with available measurements. The initial CFD scatter space is depicted by the green bar, serving as a basis for algorithm regularization. In the diagram, experimental observations at plane 4 are shown in blue. The inference mean update and its associated uncertainty are depicted in red. At plane 4, the DA method aligns with the experimental observation. At plane 2, a validation experimental measurement (not used in the DA method) falls within the scope of the DA estimation.

These examples collectively underscore the potential of DA techniques in bridging the gap between numerical simulations and real-world observations within the complexities of various flow scenarios. The ability to integrate multiple data sources effectively allows for improved accuracy, enabling applications across diverse fields.

However, while inverse problems provide a valuable framework for reconciling computational predictions with observed data, they primarily focus on refining the simulation model to align with real-world observations. In doing so, they may not inherently address the second challenge, the limited spatial and temporal coverage in experimental

data. This second challenge is precisely what the methodology proposed in this work aims to tackle using an hybrid data-driven approach. By transitioning from inverse problems to regression problems, attention is shifted to harness the power of regression-based data fusion algorithms, to merge information from different data sources and develop multi-fidelity models capable of interpolating and extrapolating finite data points. This approach is instrumental in enhancing experimental fluid mechanics methods while simultaneously reducing the reliance on extensive instrumentation and minimizing testing time.

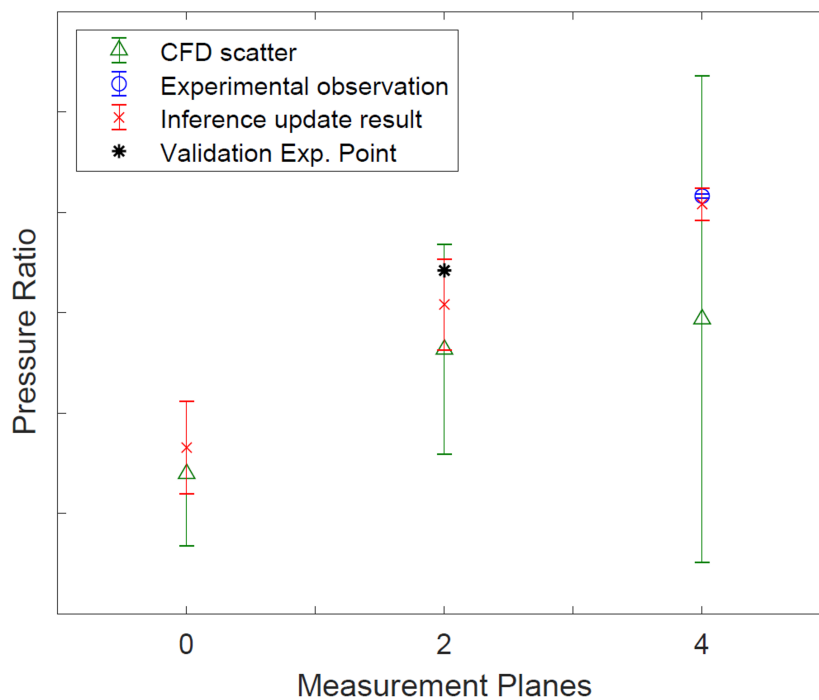


Figure 2.5. Adapted from CRUZ et al., 2021. Pressure ratio across different measurement planes. DA method matches the experimental observation at the assimilation plane. DA accurately estimates the validation experimental measurement at a different plane where the DA process was not directly influenced by the data.

2.2.2. Regression Problems

Addressing the second challenge of dealing with limited spatial and temporal coverage in experimental data means shifting from inverse problems to regression problems. This challenge is the focus of the present work, which aims to develop a hybrid data-driven methodology to tackle heavy instrumentation use and extensive testing times.

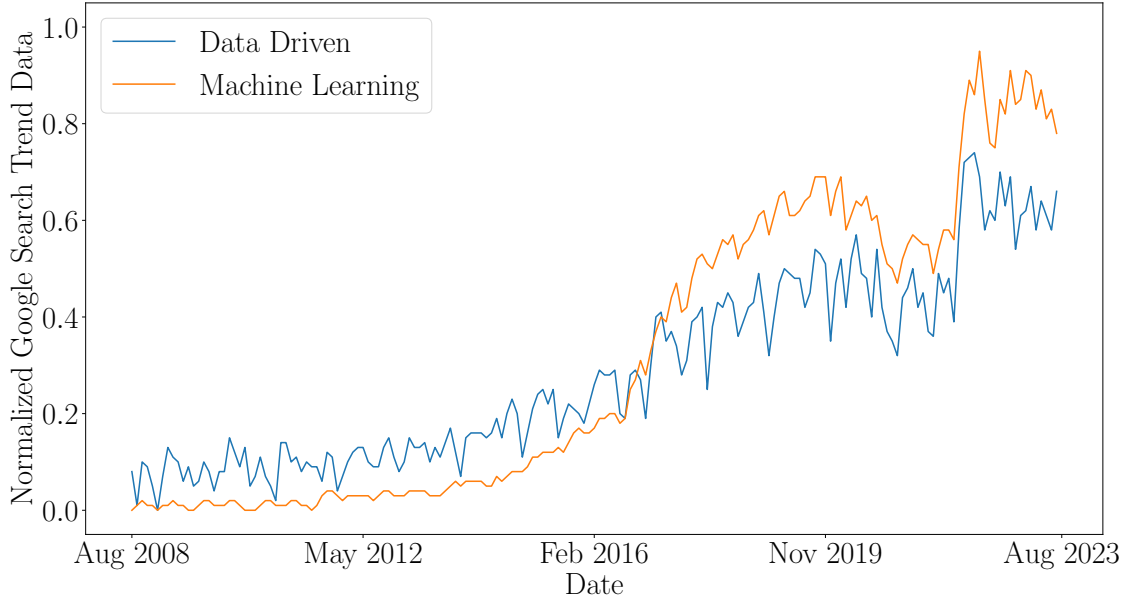


Figure 2.6. Normalized Google search trends evolution over time for the keywords 'machine learning' and 'data driven'. Both keywords follow a similar trend.

Mathematically, a regression problem involves uncovering patterns, relationships, or trends within a dataset. The primary objective of a regression model is to establish a functional mapping between input variables and output variables, enabling the prediction of response values for new or unobserved input data. The general formulation of a regression problem can be expressed as follows:

$$\mathbf{y} = f(\mathbf{X}) + \eta \quad (2.4)$$

where \mathbf{y} represents the vector of observed output variables with a random error or noise associated η . \mathbf{X} is the matrix of input variables and $f(\mathbf{X})$ denotes the underlying regression function that maps inputs to outputs.

In the context of this work, the available experimental flow field measurements \mathbf{y} at measurement locations \mathbf{X} are used to build a regression model $f(\mathbf{X})$, enabling the estimation of unmeasured flow field data points \mathbf{y}_* at locations where probe measurements are unavailable \mathbf{X}_* .

The choice of regression model $f(\mathbf{X})$ depends on the nature of the data, the relationships between variables, and previous knowledge of the dataset. Techniques range from simple linear regression to more complex methods like polynomial regression.

Recent advancements in machine learning (ML) have brought forth advanced re-

gression techniques capable of handling complex data. This is supported by Figure 2.6, which presents the normalized evolution of Google search trends for the keywords "machine learning" and "data-driven" over time. Both keywords follow a similar search trend, highlighting their correlation. The literature trends are clearly noticeable. Starting in February 2016, there has been a continuous increase in interest in machine learning that remains high to the present day.

Algorithms such as support vector regression, neural networks, and Gaussian processes offer the capacity to capture intricate relationships beyond the capabilities of traditional regression methods.

The surge of ML in fluid mechanics has not gone unnoticed, and recently, BRUNTON et al., 2020 reviewed the opportunities of ML methods in the field. One of the points highlighted is how common fluid mechanics tasks, such as reduced-order modelling, shape optimization, or flow control, can be posed as regression problems, with PINO et al., 2023 comparing different ML methods for active flow control. In parallel, one of the questions left open is how to hybridize data-driven approaches with flow physics.

Addressing this challenge, KARNIADAKIS et al., 2021 review physics-informed ML methods and explore the field of physics-informed neural networks (PINNs). PINNs combine neural networks with underlying physical principles by integrating data with partial differential equations in a loss function that is then minimized by the neural network, similar to how residuals are minimized in CFD simulations.

An interesting example is the work of CAI et al., 2021. The test case involves a 3D flow over an espresso cup measured with tomographic Schlieren imaging. They propose a PINN to leverage this data and exploit physics to predict the velocity and pressure fields from the measured temperature field. The results are shown in Figure 2.7. The reconstructed velocity and pressure fields were validated against an independent PIV experiment, and 2D plane cuts show an agreement of flow patterns.

While neural networks and other classic ML methods are good choices for building regression models $f(\mathbf{X})$, they often require a substantial amount of data to be effective. This poses a challenge in experimental fluid mechanics, where acquiring extensive datasets can be particularly demanding. Collecting high-quality data in complex flow scenarios can be costly, time-consuming, and sometimes impractical, especially when attempting to minimize instrumentation efforts and testing times, as discussed previously.

In addition, in alignment with the objectives outlined for this work, it is essential to emphasize the fundamental importance of quantifying uncertainty. Traditional regression techniques and even many ML methods may struggle to provide robust estimations




Figure 2.7. Measured temperature field data from tomographic Schlieren integrated with a PINN to exploit physics, inferring velocity and pressure fields. Adapted from CAI et al., 2021; KARNIADAKIS et al., 2021.

of uncertainty.

To address these challenges, attention shifts to Bayesian methods, specifically Gaussian processes (GPs). GPs can be effectively used in a regression formulation even with limited available data, giving a robust uncertainty quantification (UQ) due to their probabilistic nature within the Bayesian inference framework. In the following section, the Bayesian inference framework will be explored before delving into the utilization of GPs in developing the hybrid multi-fidelity experimental methodology.

2.3. Bayesian Inference: Tackling Uncertainty

Motivated by the need of robust UQ and given the primary objective of this work, of developing a hybrid multi-fidelity experimental methodology that deals with limited data availability, Bayesian inference theory becomes a fundamental tool.

Bayesian inference offers a probabilistic mathematical framework for UQ, treating both model parameters and predictions as probability distributions rather than fixed real numbers, thus allowing a rigorous uncertainty propagation, a key point when dealing with experimental data. It is generally introduced here through the Bayes theorem:

$$p(\boldsymbol{\theta}|\mathcal{D}) = \frac{p(\mathcal{D}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathcal{D})} \quad (2.5)$$

where $\boldsymbol{\theta}$ is a generic vector of quantities of interest, defined below depending on the application, and \mathcal{D} is the observed data used for the analysis. $p(\boldsymbol{\theta})$ is the *prior* probability distribution and represents the assumptions on $\boldsymbol{\theta}$ before taking into account the data. The conditional probability $p(\mathcal{D}|\boldsymbol{\theta})$ is called the *likelihood function* and it expresses how probable the data is given the possible values of $\boldsymbol{\theta}$. In practical terms, it will compare the observed data \mathcal{D} with a non-linear model prediction that depends on $\boldsymbol{\theta}$. The term $p(\mathcal{D})$ is a normalization constant that can usually be dropped in practical applications since it does not depend on $\boldsymbol{\theta}$ and equation 2.5 becomes:

$$p(\boldsymbol{\theta}|\mathcal{D}) \propto p(\mathcal{D}|\boldsymbol{\theta})p(\boldsymbol{\theta}) \quad (2.6)$$

Hence, through the computation of the *posterior* probability distribution $p(\boldsymbol{\theta}|\mathcal{D})$, the Bayesian framework allows inferences to be made about uncertain quantities of interest $\boldsymbol{\theta}$ based of observed data \mathcal{D} . Bayesian inference is analogous to the research methodology WIKLE and BERLINER, 2007. To prior research, one adds observations/experiments and infers new conclusions.

Figure 2.8 illustrates a schematic of the Bayes theorem, represented in equation 2.5, to infer the posterior distribution of a single variable θ . An initial belief of the probability distribution of θ is depicted in blue. It has a large standard deviation representing the high uncertainty associated with the variable. A likelihood function $p(\mathcal{D}|\theta)$, in green, expresses the joint probability of the observed data \mathcal{D} given the prior distribution of θ , i.e. as a measure of how well a model, based on θ , explains the observed data. The Bayesian update estimates the posterior probability distribution $p(\theta|\mathcal{D})$, in red. The inferred distribution, unlike the original prior, shows a smaller standard deviation and a shifted mean, proving a more accurate estimation of θ based on a foundational statistical theorem.

In the previous sections, various works were mentioned where Bayesian inference works as the core foundational theory. In the context of DA inverse problems 2.2.1, a direct parallel can be drawn between equation 2.3 and 2.6. The input parameters $\boldsymbol{\theta}$ are given a prior probability distribution, which is then updated based on experimental observations (new data with associated uncertainty η), and a posterior distribution is computed through the DA algorithm.

Selecting an appropriate prior probability distribution is an important step in the Bayesian process, since it allows for the addition of physical knowledge to the model.




Figure 2.8. Schematic example of Bayesian inference. Bayesian inference refines the initial belief (prior distribution in blue) about a variable (θ) when new observed data (\mathcal{D}) is considered. The updated belief (posterior distribution in red) provides a more accurate estimate of the variable.

For example, if the input parameter represents a temperature value, it can be assumed that the value will always be higher than 0 K. Additionally, having an 'informed' prior, characterized by physically constrained prior distributions, serves as a form of regularization, particularly in ill-posed problems ZHANG et al., 2020.

In the context of regression, Bayes theorem is often used in Bayesian linear regression. Here, the regression model $f(\mathbf{X})$ in equation 2.4 is a linear combination of weights and observations. By assigning specific prior probability distributions to these weights, the posterior distribution is analytical tractable and uncertainty can be directly propagated when making predictions. A mathematical overview of Bayesian based methods is given by BISHOP, 2006.

Bayesian statistics can also be extended to non-linear machine learning methods, such as neural networks, giving rise to Bayesian neural networks, introduced by BLUNDELL et al., 2015. Similar to the case of Bayesian linear regression, probability distributions are assigned to the network's weights and biases, enabling the propagation of uncertainty in predictions. It is worth noting that this approach, while robust, often comes with a

significantly higher computational cost compared to standard neural networks, which are already computationally intensive due to the large volume of data they require.

Another Bayesian approach to regression problems is Gaussian process regression (GPR). GPR leverages GPs, incorporating probabilistic distributions directly into the regression model $f(\mathbf{X})$. While GPR can be computationally expensive, it excels when dealing with limited data, offsetting the higher computational cost.

In the context of this work, GPs have been chosen as the Bayesian algorithm for the hybrid multi-fidelity experimental methodology. This choice is motivated by the alignment of ML regression trends with the robust statistical framework for UQ. The following section provides an introduction to GPs modelling, starting with the fundamental approach and leading to the multi-fidelity modelling approach adopted in this work.

2.4. Gaussian Process Modelling

Before introducing the multi-fidelity Gaussian Process (MFGP) approach that is proposed for the hybrid experimental methodology of this work, it is necessary to take a step back and introduce a general overview GPs.

GPs are a non-parametric method that constructs “Gaussian distributions over functions” RASMUSSEN and WILLIAMS, 2005 rather than in a fixed number parameters. This unique characteristic makes GPs a powerful probabilistic tool, especially for regression problems, as formulated in equation 2.4, predict unknown data \mathbf{y}_* from observed data \mathbf{y} , which includes some level of measurement error or uncertainty.

Fundamentally, GPs leverage the special properties of Gaussian distributions to extend joint multivariate Gaussian distributions to an infinite number of points. This extension allows to model a continuous function $f(\mathbf{x})$ over the input domain \mathbf{x} . To understand how GPs construct distributions over functions, and assuming a basic statistical knowledge, the basics of multivariate Gaussian distributions are introduced.

A multivariate Gaussian distribution, also known as a multivariate normal distribution, is a probability distribution that describes the joint distribution of multiple random variables, which are themselves Gaussian distributed. In a multivariate context, a Gaussian distribution is fully defined by a mean vector, denoted $\boldsymbol{\mu}$, and a covariance matrix, here initially denoted as $\boldsymbol{\Sigma}$.

In mathematical terms, given a dataset matrix $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n] \in \mathbb{R}^{k \times n}$ a multivariate Gaussian distribution of a k -dimensional random vector $\mathbf{x} = [x_1, x_2, \dots, x_k]^T$ is

represented as:

$$\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \quad (2.7)$$

where $\boldsymbol{\mu}$ is a k -dimensional vector containing the means of the individual random variables and $\boldsymbol{\Sigma}$ is a matrix ($k \times k$ size) that specifies how these variables are correlated. The covariance matrix captures both the variances (diagonal elements) and the covariances (off-diagonal elements) between variables.

With a known mean and a positive definite covariance, the probability density function of the multivariate Gaussian distribution is given by:

$$p(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{\sqrt{(2\pi)^k |\boldsymbol{\Sigma}|}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right) \quad (2.8)$$

Gaussian distributions have some fundamental properties that allow for GPs to analytically solve Bayes theorem 2.6.

One of the most important properties is that Gaussian distributions maintain their Gaussian nature when subjected to affine transformations. Mathematically, if \mathbf{x} follows a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, then applying an affine transformation $\mathbf{L}\mathbf{x} + \mathbf{b}$ to \mathbf{x} results in another Gaussian distribution:

$$\mathbf{L}\mathbf{x} + \mathbf{b} \sim \mathcal{N}(\mathbf{L}\boldsymbol{\mu} + \mathbf{b}, \mathbf{L}\boldsymbol{\Sigma}\mathbf{L}^T) \quad (2.9)$$

The other two relevant Gaussian properties for GPs are marginalization and conditioning operations.

A marginal distribution represents the probability distribution of a subset of the random variables while marginalizing ('referencing') the original joint distribution. Given a joint Gaussian distribution over variables $\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2]$, the marginal distribution of \mathbf{x}_1 can be obtained as follows:

$$p(\mathbf{x}_1) = \int p(\mathbf{x}_1, \mathbf{x}_2) d\mathbf{x}_2 \sim \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_{11}) \quad (2.10)$$

Figure 2.9 shows the pdf of the joint 2-dimensional Gaussian distribution and the marginal distributions, $p(\mathbf{x}_1)$ and $p(\mathbf{x}_2)$ that are obtained with the marginalization property of equation 2.10.

Similarly, conditioning on observed data \mathbf{x}_1 to update our beliefs about the remaining variables \mathbf{x}_2 , yielding a conditional distribution that is directly comparable to the Bayes theorem in equation 2.5:

$$p(\mathbf{x}_2 \mid \mathbf{x}_1) = \frac{p(\mathbf{x}_1 \mid \mathbf{x}_2)p(\mathbf{x}_2)}{p(\mathbf{x}_1)} \quad (2.11)$$

The formulation with just Gaussian distributions allows the posterior to be analytically tractable as it will be presented below.

These properties enable GPs to make predictions and infer uncertainty from limited observed data. When some sparse data is available \mathbf{x}_1 , what is the expectation about other unseen data \mathbf{x}_2 , assuming that they come from the same joint distribution.

With the foundation and notation of multivariate Gaussian distributions established, the focus is now on the application of these principles in Gaussian Processes for function modelling.




Figure 2.9. 2-dimensional multivariate Gaussian distribution. The joint probability distribution $p(\mathbf{x}_1, \mathbf{x}_2)$ is denoted at the center. The marginal distributions $p(\mathbf{x}_1)$ and $p(\mathbf{x}_2)$ are shown as projections.

2.4.1. Introduction to Gaussian Process

A Gaussian process defines a multivariate Gaussian prior distribution over functions for every point in the input space \mathbf{X} :

$$p(f \mid \mathbf{X}) \sim \mathcal{N}(f \mid \boldsymbol{\mu}, \mathbf{K}) \quad (2.12)$$

where $f(\mathbf{X})$ is a multivariate random variable $f(\mathbf{X}) = [f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_N)]$ with a mean function $\boldsymbol{\mu}$ and a covariance matrix $\mathbf{K}(\mathbf{x}_i, \mathbf{x}_j)$ that is defined with a positive definite kernel function. Due to the flexibility of GPs in modelling, the mean is assumed to be zero $\boldsymbol{\mu} = 0$ and the GP distribution over $f(\mathbf{X})$ is fully defined by \mathbf{K} . Therefore, the choice of the kernel function \mathbf{K} is a key modelling decision in the GP framework.

Multiple kernel function families exist in the literature for different function modelling behaviour (e.g. smoothness, stationary) making this choice a prior input of knowledge of the problem. In the present work, unless specified otherwise, the base kernel function selected is the stationary anisotropic Matérn-5/2 kernel, which is given by:

$$\mathbf{K}(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \left(1 + \frac{\sqrt{5 (\mathbf{x}_i, \mathbf{x}_j)^T (\mathbf{x}_i, \mathbf{x}_j)}}{\ell} + \frac{5 (\mathbf{x}_i, \mathbf{x}_j)^T (\mathbf{x}_i, \mathbf{x}_j)}{3\ell^2} \right) \exp \left(-\frac{\sqrt{5 (\mathbf{x}_i, \mathbf{x}_j)^T (\mathbf{x}_i, \mathbf{x}_j)}}{\ell} \right) \quad (2.13)$$

where ℓ is the length scales for a single input dimension and control the smoothness of the function modelling and σ_f acts as a scaling factor. These are usually called hyperparameters $\boldsymbol{\theta}$ and ensuring their correct values play a pivotal role in GPs. The length scales ℓ controls the smoothness of the function by directly quantifying how much a single data point influences the space around it. The scaling factor σ_f controls the vertical scale of the function around its mean, influencing the uncertainty of the model.

Given the above assumptions and available noisy observations $\mathbf{y}_i = f(\mathbf{x}_i) + \boldsymbol{\eta}_i$, where the noise η follows an independent, identically distributed Gaussian distribution with zero mean and fixed variance $\eta \sim \mathcal{N}(\mathbf{0}, \sigma_y^2 \mathbf{I})$, the log marginal likelihood can be derived as:

$$\log p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\theta}) = -\frac{1}{2} \log (\det (\mathbf{K} + \sigma_y^2 \mathbf{I})) - \frac{1}{2} (\mathbf{y} - \boldsymbol{\mu})^\top (\mathbf{K} + \sigma_y^2 \mathbf{I})^{-1} (\mathbf{y} - \boldsymbol{\mu}) - \frac{n}{2} \log(2\pi) \quad (2.14)$$

where the first term is a model complexity term, defined by the selection of the kernel function. The second term is a likelihood data-fit term showing the Bayesian weighting

of the prior with the observed data. The third term is a constant term. Based on the available data, the log marginal likelihood can be maximized to obtain the optimal hyperparameters mentioned above.

The inherent noise of the available data σ_y can be fixed to a known value, but it can also be taken as an hyperparameter to infer along with the kernel parameters. In practical terms, it is not uncommon to treat the noise parameter σ_y as a hyperparameter in the GP model. This choice allows the model to determine the optimal level of noise during hyperparameter training. It enhances model flexibility by enabling adaptation to varying noise levels and it provides valuable uncertainty information of what the model thinks of the data. Moreover, from a Bayesian perspective, a prior beliefs about noise can be updated based on observed data. However, it IS essential to note that treating σ_y as a hyperparameter increases the complexity of the model and having a poor initial prior on it can lead to incorrect results with GPs models explaining the entire available data just with noise.

Predictions of the unknown states \mathbf{y}_* can be obtained by conditioning their input locations \mathbf{X}_* on the prior distribution of equation 2.12, giving, a predictive posterior distribution:

$$p(\mathbf{y}_* | \mathbf{X}_*, \mathbf{X}, \mathbf{y}) = \mathcal{N}(\mathbf{y}_* | \boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*) \quad (2.15)$$

To reach the posterior solution, a joint distribution of the observed values \mathbf{y} and the unknown states \mathbf{y}_* , which is itself a Gaussian can be written as:

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}_* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \mathbf{K} & \mathbf{K}_*^T \\ \mathbf{K}_* & \mathbf{K}_{**} \end{bmatrix}\right) \quad (2.16)$$

where $\mathbf{K}_* = \mathbf{K}(\mathbf{x}_*, \mathbf{x})$ and $\mathbf{K}_{**} = \mathbf{K}(\mathbf{x}_*, \mathbf{x}_*)$.

Using the Gaussian properties of marginalization and conditioning from equations 2.10 and 2.11, respectively, the predictive distribution is given by:

$$\boldsymbol{\mu}_* = \mathbf{K}_*(\mathbf{K} + \sigma_y^2\mathbf{I})^{-1}\mathbf{y} \quad (2.17)$$

$\boldsymbol{\mu}_*$ represents the predicted mean values for new input locations \mathbf{X}_* based on the observed data \mathbf{y} , and it is influenced by the covariance between the observed data and the new data locations, and the observed data itself.

$$\boldsymbol{\Sigma}_* = \mathbf{K}_{**} - \mathbf{K}_*(\mathbf{K} + \sigma_y^2\mathbf{I})^{-1}\mathbf{K}_*^T \quad (2.18)$$

where $\boldsymbol{\Sigma}_*$ is the covariance matrix associated with the predicted values, providing information about the uncertainty or variability in predictions with the standard deviation

being computed directly from the diagonal. It depends on the covariances between the new data locations and between the observed and new data, adjusted by the inherent noise of the observations. This showcases the non-parametric modelling property of GPs, which do not summarize data with parameters but instead rely on the available data for predictions at the cost of a higher computational cost.

Having established GPs for function modelling, the following subsection extends this framework to address scenarios involving multi-fidelity data. In the context of this work, multi-fidelity Gaussian processes become valuable tools for leveraging data from various sources and fidelities to enhance modelling capabilities. This extension is particularly significant to develop a hybrid multi-fidelity experimental methodology.

2.4.2. Multi-Fidelity Gaussian Process

The multi-fidelity modelling literature originates from a computational model management perspective, PEHERSTORFER et al., 2018 divides the multi-fidelity approaches in three categories: adaptation, fusion and filtering. In this work, multi-fidelity Gaussian processes (MFGP) are used in a data fusion context, where low- and high-fidelity 'models' are available and their output information is combined for an improved model.

Multi-fidelity modelling surges from the need to couple complex numerical codes with different computational costs either to speed up optimization or for numerical uncertainty quantification, through surrogate models FERNÁNDEZ-GODINO, 2016. In their review work, only 5% of the multi-fidelity literature was exploiting simulations with experimental data, with this being evident also in fluid mechanics. This gap of literature falls in the context of this work, that will leverage experimental data as 'high-fidelity data' and thus develop a hybrid measurement technique.

The review of multi-fidelity modelling approaches based on GPs, conducted by BREVAULT et al., 2020, presents different algorithms to couple information between varying fidelity levels of data. The different algorithms have two main distinctions. Figure 2.10 illustrates the classification of the literature most common approaches within these main distinctions: the symmetry/asymmetry of the fidelity levels and their linear/non-linear relationship.

Symmetry in this context refers to the hierarchical weight of the different fidelities. In symmetrical approaches, all levels of fidelity are treated equally, and the models aim to leverage information through the outputs correlation. This approach is often referred to as co-kriging. While it simplifies the modelling process and can be effective when all fidelities are equally reliable, it becomes disadvantageous when dealing with

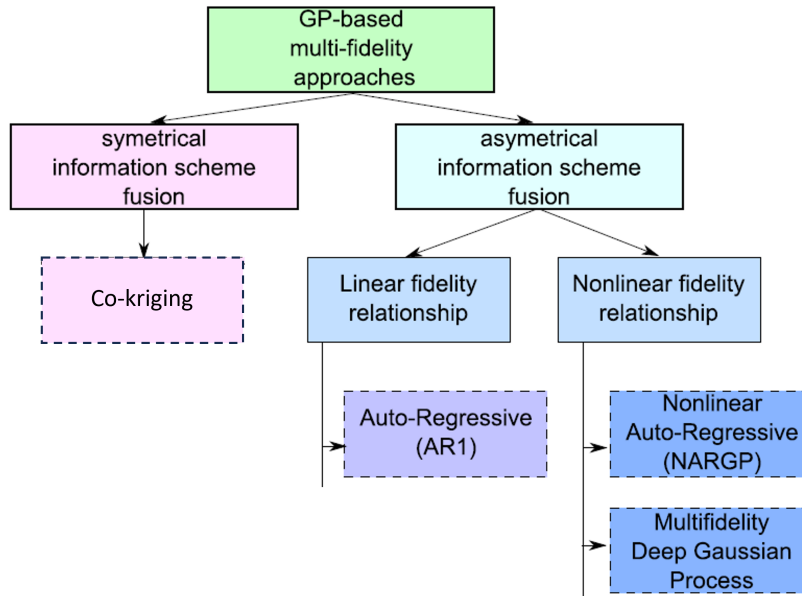


Figure 2.10. Classification of multi-fidelity GP approaches based on symmetric/asymmetric data fusion scheme and linear/non-linear relationship between different fidelity levels. Adapted from BREVAULT et al., 2020.

asymmetrical information.

In cases where certain fidelities are more reliable or informative than others, treating them symmetrically can lead to suboptimal results. For instance, when a high-fidelity experimental dataset and a low-fidelity numerical simulation are available, asymmetrical methods give more weight to the high-fidelity data, while co-kriging may not fully exploit the valuable information provided by the high-fidelity source.

The linearity of relationships between fidelity levels is another modelling choice. Classical multi-fidelity methods, originally introduced by KENNEDY and O’HAGAN, 2000, assume linear dependencies between these levels. More recent data heavy approaches, such as Non-linear Auto-Regressive multi-fidelity Gaussian Process (NARGP) PERDIKARIS et al., 2017 and Multi-Fidelity Deep Gaussian Process (MF-DGP) DAMIANOU and LAWRENCE, 2013 use a non-linear fidelity coupling that is able to capture more complex interactions at the cost of requiring more data and higher computational costs.

In this work, the proposed methodology leverages MFGP. Specifically, it relies on the use an auto-regressive linear method (AR1) data fusion scheme across different fidelity levels. It offers a simple but effective framework for modelling the relationships between different fidelity levels. Importantly, it retains the core structure of GPs, enabling to make predictions while quantifying the associated uncertainty.

The AR1 formulation assumes a linear dependency between each fidelity model. In a general formulation, it assigns a GP prior to each fidelity model t , where the higher fidelity model f_t is expressed as a function of the lower fidelity f_{t-1} , multiplied by a scaling factor $\rho_{t-1}(\mathbf{x})$ plus a bias function $\delta_t(\mathbf{x})$, which is itself a GP $\delta_t(\mathbf{x}) \sim \mathcal{N}(\mu_\delta, K_t)$.

$$f_t(\mathbf{x}) = \rho_{t-1}(\mathbf{x})f_{t-1}(\mathbf{x}) + \delta_t(\mathbf{x}) \quad (2.19)$$

The scaling factor $\rho_{t-1}(\mathbf{x})$ weights the cross-correlation between fidelity levels and it is considered an hyperparameter to be learned in the inference process and is usually assumed to be a constant. The choice of whether $\rho_{t-1}(\mathbf{x})$ is treated as a constant or varies spatially can depend on the specific problem and the available data, but in practice, varying $\rho_{t-1}(\mathbf{x})$ leads to an increased complexity in the inference process and increased computational cost and thus is treated as a constant and added to the kernel vector of hyperparameters $\boldsymbol{\theta}$ to infer.

Practically, $\rho_{t-1}(\mathbf{x})$ quantifies how much influence the lower-fidelity model $f_{t-1}(\mathbf{x})$ has on the higher-fidelity model $f_t(\mathbf{x})$. When $\rho_{t-1}(\mathbf{x})$ is close to 1, it indicates a strong linear relationship, implying that changes in the lower-fidelity model strongly affect the higher-fidelity model. Conversely, when $\rho_{t-1}(\mathbf{x})$ is smaller, it suggests a weaker relationship, with the higher-fidelity model being less sensitive to changes in the lower-fidelity model.

The above formulation ensures that at each fidelity level t , the conditional distribution of the GP $f_t(\mathbf{x})$ is influenced by $f_{t-1}(\mathbf{x})$, with both the scaling factor and the bias function contributing to the overall fidelity transition. The MFGP AR1 formulation of equation 2.19 is thus derived from the Markov property:

$$\text{Cov}\{f_t(\mathbf{x}), f_{t-1}(\mathbf{x}') \mid f_{t-1}(\mathbf{x})\} = 0, \quad \forall \mathbf{x} \neq \mathbf{x}' \quad (2.20)$$

which means that assuming $f_{t-1}(\mathbf{x})$ is known, no more can be learned about $f_t(\mathbf{x})$ from any other lower fidelity model output $f_{t-1}(\mathbf{x}')$, for $\mathbf{x} \neq \mathbf{x}'$.

In the case where only 2 levels of fidelity are available, which is the default case in this work, and assuming a constant scaling factor ρ , the formulation of equation 2.19 simplifies to:

$$f_{high}(\mathbf{x}) = \rho f_{low}(\mathbf{x}) + \delta(\mathbf{x}) \quad (2.21)$$

where $f_{high}(\mathbf{x})$ and $f_{low}(\mathbf{x})$ represent the high- low-fidelity models, respectively. The prior GP model of the formulation of equation 2.21 can be expressed as:

$$\begin{bmatrix} f_{low}(\mathbf{x}) \\ f_{high}(\mathbf{x}) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \boldsymbol{\mu}_{low} \\ \boldsymbol{\mu}_{high} \end{bmatrix}, \begin{bmatrix} \mathbf{K}_{low} & \rho\mathbf{K}_{low} \\ \rho\mathbf{K}_{low} & \rho^2\mathbf{K}_{low} + \mathbf{K}_{high} \end{bmatrix} \right) \quad (2.22)$$

where $\boldsymbol{\mu}_{low}$ and $\boldsymbol{\mu}_{high}$ are the mean functions. \mathbf{K}_{low} is the covariance matrix associated with the low-fidelity data, $\rho\mathbf{K}_{low}$ is the cross correlation matrix between fidelities and $\rho^2\mathbf{K}_{low} + \mathbf{K}_{high}$ is the covariance matrix associated with the high-fidelity.

There are two main numerical schemes for inferring the GPs posterior: a fully coupled formulation by KENNEDY and O'HAGAN, 2000 and a recursive approach developed by LE GRATIET and GARNIER, 2014 with an thorough mathematical derivation found in LE-GRATIET, 2013.

The recursive approach of LE GRATIET and GARNIER, 2014 decouples the inference of the fidelities into standard GP inference by assuming a nested data structure $\mathcal{D}_{high} \subseteq \mathcal{D}_{low}$. The key step of the formulation is to infer the posterior of the low-fidelity model $f_{low}^{\sim}(\mathbf{x})$ independently and replace it for its prior in equation 2.21. With this formulation, the multi-fidelity GP posterior distribution for the high fidelity $p(f_{high} | \mathbf{X}_{high}, \mathbf{y}_{high}, f_{low}^{\sim})$ is fully defined and predictions are made with following mean and covariance:

$$\mu_{high,*}(\mathbf{x}_*) = \rho\mu_{low}(\mathbf{x}_*) + \mu_{\delta} + \mathbf{K}_{high,*} (\mathbf{K}_{high} + \sigma_{y,high}^2 \mathbf{I})^{-1} [\mathbf{y}_{high} - \rho\mu_{low}(\mathbf{x}_{high}) - \mu_{\delta}] \quad (2.23)$$

$$\boldsymbol{\Sigma}_{high,*}(\mathbf{x}_*) = \rho^2 \boldsymbol{\Sigma}_{low}(\mathbf{x}_*) + \mathbf{K}_{high,**} - \mathbf{K}_{high,*} (\mathbf{K}_{high} + \sigma_{y,high}^2 \mathbf{I})^{-1} \mathbf{K}_{high,*}^T \quad (2.24)$$

Built upon this MFGP mathematical framework, BABAEE et al., 2016 apply MFGP to revisit a classic problem in fluid dynamics: the flow past a cylinder, specifically studying thermal mixed convection. By combining high-fidelity DNS data with a low-fidelity empirical correlation, they constructed a MFGP model capable of predicting the Nusselt number as a function of Reynolds and natural-convection directions. Leveraging the the multi-fidelity framework, they explored a comprehensive range of conditions supported by sparse DNS data to explore the underlying physics of the problem and propose a new correlation based on a physics based correction of the Reynolds number definition that better matches the model.

The multi-fidelity Gaussian processes framework presented here is proven powerful to the development of a hybrid measurement technique that, according to the objectives

set out in this work, will couple experimental data with numerical models.

However, the tools presented so far are limited to already having data, meaning that an experimental campaign could not be planned in a proactive manner just with these tools. To later address this objective, in the next subsection, Sparse Variational Gaussian Processes are introduced in a general way.

2.4.3. Stochastic Variational Gaussian Process

The GPs discussed in the preceding sections provide a robust framework for modelling complex relationships in data, making predictions, and quantifying uncertainties. However, in many real-world applications, obtaining an extensive dataset may be impractical and expensive, or simply infeasible. Moreover, traditional GP models struggle to handle such large datasets efficiently due to their cubic computational complexity with respect to the number of data points. To address these limitations, HENSMAN et al., 2013 developed Sparse Variational Gaussian Process (SVGP).

In layman's terms, the core idea behind SVGP is to select a smaller subset of "inducing points" from the dataset and then build a GP model based on this reduced set. These inducing points act as a representative sample of the data.

Formally, a SVGP model is a scalable approximation to the GP that is designed to account for large datasets by employing a sparse subset of pseudo-inputs, referred to as inducing points \mathbf{Z} .

The SVGP employs a variational distribution $q(f | \mathbf{X})$ which is more tractable than the original distribution, but is designed to be a close approximation. This allows for a significant decrease in computational complexity, with the cost now scaling linearly with the number of data points n rather than cubically as in the full GP. The variational distribution is defined as follows:

$$q(f | \mathbf{X}) = \int p(f | \mathbf{u}, \mathbf{X})q(\mathbf{u})d\mathbf{u} \quad (2.25)$$

where \mathbf{u} represents the function values at the pseudo-inputs, and $q(\mathbf{u})$ is the variational distribution over these function values. $q(\mathbf{u})$ is chosen as a Gaussian:

$$q(\mathbf{u}) = \mathcal{N}(\mathbf{u} | \mathbf{m}, \mathbf{S}) \quad (2.26)$$

where \mathbf{m} and \mathbf{S} are variational parameters. The pseudo-inputs and variational parameters are considered hyperparameters that are jointly optimized to minimize the Kullback-Leibler divergence between the variational distribution $q(f | \mathbf{X})$ and the true

posterior $p(f \mid \mathbf{X}, \mathbf{y})$.

The variational evidence lower bound (ELBO) is equivalent of the log marginal likelihood in the SVGP framework and can be written as:

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\lambda}, \mathbf{Z}) = \mathbb{E}_{q(f|\mathbf{X})} [\log p(\mathbf{y} \mid f, \mathbf{X})] - \text{KL} [q(\mathbf{u}) \parallel p(\mathbf{u})] \quad (2.27)$$

where $\boldsymbol{\lambda} = \mathbf{m}, \mathbf{S}$ are the variational parameters, and KL denotes the Kullback-Leibler divergence. The expectation is taken with respect to the variational distribution $q(f \mid \mathbf{X})$. Similarly to classic GPs, the ELBO is maximized to obtain the optimal hyperparameters. This is where the 'stochastic' name comes from, as this optimization problem is solved with stochastic gradient descent (SGD) optimization. In the context of SVGP, SGD is employed to jointly optimize the kernel parameters, variational distribution parameters (\mathbf{m} and \mathbf{S}), and the locations of the inducing points (\mathbf{Z}).

SGD is a powerful optimization technique that works by iteratively updating the hyperparameters in the direction that minimizes a cost function in this case the ELBO. Unlike traditional gradient descent, which computes the gradient using the entire dataset (batch gradient descent), SGD operates on random mini-batches of data. This randomness introduces noise into the optimization process but has the advantage of being computationally efficient and capable of handling large datasets.

Finally, the predictive posterior distribution in SVGP is obtained by conditioning on the variational distribution:

$$p(\mathbf{y}_* \mid \mathbf{X}_*, \mathbf{X}, \mathbf{y}, \mathbf{Z}) = \mathcal{N}(\mathbf{y} \mid \boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*) \quad (2.28)$$

The predictive mean $\boldsymbol{\mu}_*$ and covariance $\boldsymbol{\Sigma}_*$ can be calculated similarly to the full GP case, but now involving the inducing points and covariance matrices \mathbf{K}_Z and $\mathbf{K}_{\mathbf{XZ}}$ that define an approximated covariance matrix and the cross-covariance matrix between the training data \mathbf{X} and the inducing points \mathbf{Z} .

In this concluding section, the mathematical framework and modelling approaches that are necessary to achieve the proposed goal of developing an hybrid experimental methodology. MFGP will perform the integration of experimental and numerical data to evaluate flow characteristics, effectively reducing the need for extensive experimental testing. Additionally, SVGP will be instrumental within a novel DoE framework, enabling the efficient planning of data sampling locations.

The next chapter will introduce the methodology itself, providing a comprehensive insight into its elements, implementation strategies, and the practical application of these mathematical tools. This chapter will illustrate how these tools are employed to

effectively address the research objectives of developing a robust hybrid experimental methodology that leverages different data sources while ensuring an accurate uncertainty quantification.

Chapter 3.

Hybrid Measurement Technique: Structure and Methodology

IN this chapter, the proposed hybrid measurement technique is presented, breaking down each building block needed for accurate flow assessment. The primary goal is to effectively merge experimental and numerical data using Gaussian Processes, as discussed earlier, while strategically planning data collection with a Design of Experiments approach. This chapter offers practical insights into implementing these components.

Overview

3.1. Introduction	36
3.2. Axial Compressor Flow Key Parameters	38
3.3. Multi-Fidelity Gaussian Process: Implementation	39
3.3.1. Kernel selection and influence	40
3.3.2. Hyperparameters optimization	42
3.3.3. Computational implementation	43
3.3.4. Implementation and example: Forrester function	44
3.4. Design of experiments: <i>A Priori</i> sampling	46
3.4.1. SVGP as DoE tool: Implementation	47
3.5. Adaptive <i>In situ</i> sampling	52
3.6. Methodology Performance Evaluation	55

3.1. Introduction

A survey of literature exploring various data-driven approaches for handling experimental data reveals the significant demand for extensive datasets in various applications. In this context, the objectives of this work align with the mathematical framework of regression problems, as discussed in data assimilation literature. Guided by current trends in machine learning for regression tasks keeping in mind the need for the mitigation of instrumentation and minimization of data acquisition times, GP modelling emerges as a pivotal tool for developing a hybrid measurement technique.

The primary goal of this work is to introduce a methodology that is designed to facilitate the performance characterization and flow assessment of turbomachinery components. The proposed approach aims to circumvent the limitations posed by instrumentation and significantly reduce testing durations.

To illustrate the methodology, Figure 3.1 provides a schematic representation. This schematic builds upon the notation established in the previous chapter and presents a visual overview of the proposed data-driven Bayesian hybrid measurement technique, which enables accurate flow assessment and uncertainty quantification while minimizing instrumentation usage and testing time.

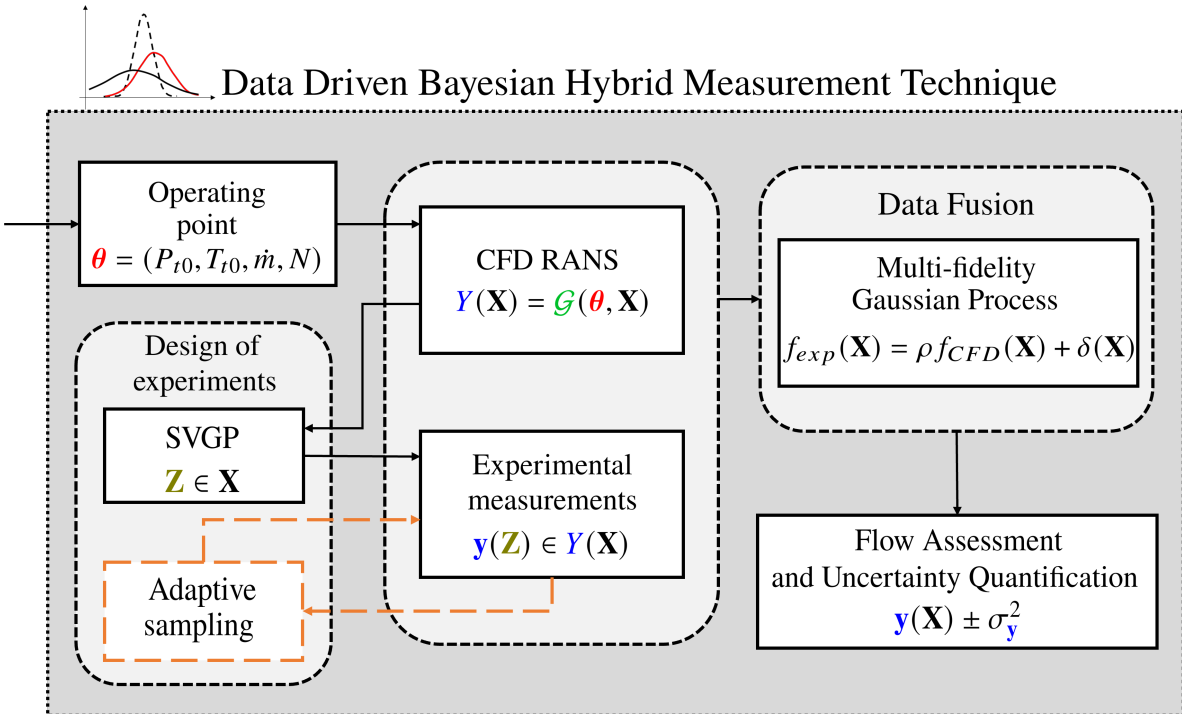


Figure 3.1. Proposed Data Driven Bayesian Hybrid Measurement Technique schematic for a complete accurate flow assessment and uncertainty quantification with a reduction in instrumentation usage and testing time.

The methodology begins with the selection of an operating point (OP) from the test matrix. From the OP conditions, a parameter vector denoted as $\boldsymbol{\theta}$ is established. This parameter encompasses the boundary conditions required to perform a CFD RANS simulation. In turbomachinery applications, these parameters typically are total inlet conditions, mass flow rate, and rotational speed, denoted respectively as $\boldsymbol{\theta} = (P_{t0}, T_{t0}, \dot{m}, N)$.

Subsequently, a CFD RANS simulation, referred to as \mathcal{G} in Figure 3.1, is conducted using the established boundary conditions for the entire turbomachinery domain under investigation. This simulation generates the desired flow field quantities across its entire spatial domain, denoted as $Y(\mathbf{X})$.

Following the CFD simulation, the flow field information serves as input for a DoE framework. Within this framework, a SVGP model is trained to *a priori* select experimental spatial locations denoted as \mathbf{Z} for sampling the flow field of interest. Based on this DoE, experimental measurements are acquired at the chosen spatial locations, yielding $\mathbf{y}(\mathbf{Z})$.

An additional innovative option within this DoE framework is the use of an adaptive sampling approach, based on the concept of Bayesian optimization, represented by the dashed orange lines in the Figure. This approach, explained in further detail later in this chapter, performs *in situ* data driven sampling. Based on a set of experimental measurements \mathbf{y}_0 , an initial GP model is trained. Subsequently, guided by this model and considering its uncertainty, the next experimental point is selected, and data is acquired (y_{m+1}). This process is iterative until specific stopping criteria are met. Following this phase, the methodology follows the same steps as with the "classic" DoE framework.

The data fusion of the CFD RANS data and experimental measurements is accomplished using a MFGP model. In this model, the highly discretized CFD RANS data and the sparsely sampled experimental measurements are treated as low- and high-fidelity data, respectively, and modelled using the autoregressive scheme as presented previously $f_{exp}(\mathbf{X}) = \rho f_{CFD}(\mathbf{X}) + \delta(\mathbf{X})$.

Once the MFGP model is trained, it enables the inference of the flow field across the entire spatial domain with an accuracy level comparable to physical experimental measurements. Finally, the Bayesian framework allows uncertainty quantification for the reconstructed flow, providing $\mathbf{y}(\mathbf{X}) \pm \sigma_{\mathbf{y}}^2$.

In the following sections, the key components of the methodology, namely the data fusion with MFGP and the design of experiments are explored. These sections provide a detailed explanation of each component, offering detailed insights into their imple-

mentation, methodologies, and practical considerations. To make informed modelling choices, it is important to first characterize the type of measurements the proposed methodology will focus on. Therefore, the next section gives a brief introduction to axial compressor flow and its essential parameters.

3.2. Axial Compressor Flow Key Parameters

This section introduces key parameters for characterizing the flow in an axial compressor. These parameters play the main role in assessing the compressor's performance, making them fundamental to the proposed methodology.

An axial compressor stage consists of two blade rows: the rotor, a rotating row of blades, give kinetic energy to the fluid, increasing its absolute velocity. Subsequently, the stator, a stationary row of blades, decelerates the flow, converting kinetic energy into static pressure. This process is graphically represented on a temperature-entropy ($T-s$) diagram, as depicted in Figure 3.2. The diagram illustrates the thermodynamic changes undergone by the fluid as it passes through the compressor stage, including variations in total pressure and temperature.

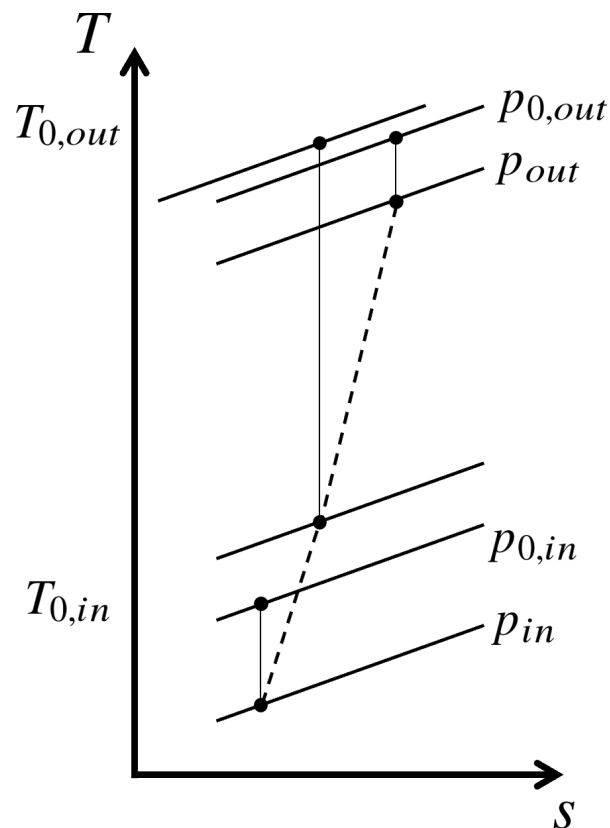


Figure 3.2. T-s diagram of an axial compressor stage.

Two key parameters, the total pressure ratio (Π) and the isentropic efficiency (ι), are central to evaluating axial compressor performance under specific operating conditions, under a fixed rotational speed (N) and mass flow rate (\dot{m}).

The total pressure ratio quantifies the change in total pressure from the compressor's inlet to its outlet. It is defined as the ratio of total pressure at the outlet ($p_{0,out}$) to that at the inlet ($p_{0,in}$), as expressed in equation 3.1.

$$\Pi = \frac{p_{0,out}}{p_{0,in}} \quad (3.1)$$

The isentropic efficiency, denoted as ι , offers insights into the compressor's effectiveness in converting kinetic energy into pressure while minimizing losses. It is determined by equation 3.2, which relates Π to the ratio of total temperatures at the outlet ($T_{0,out}$) and inlet ($T_{0,in}$). The specific heat ratio (γ) characterizes the fluid's thermodynamic behaviour.

$$\iota = \frac{\Pi^{\frac{\gamma-1}{\gamma}} - 1}{\frac{T_{0,out}}{T_{0,in}} - 1} \quad (3.2)$$

These equations underscore the critical importance of accurate measurements of total pressure and total temperature at the inlet and outlet of the axial compressor. These measurements serve as the foundation for evaluating the compressor's performance and enable the assessment of flow conditions and the efficiency of the thermodynamic process within the compressor stage.

3.3. Multi-Fidelity Gaussian Process: Implementation

Having established a comprehensive understanding of the critical parameters in axial compressor flow and the significance of total pressure and total temperature measurements in the proposed methodology, this section delves into the specific choices made to customize the MFGP model for its application. This mainly includes factors such as the problem's dimensionality and complexity.

The approach taken in MFGP implementation depends significantly on the characteristics of the specific problem at hand. In the context of applying the proposed methodology to axial compressor components, the dimensionality of the GP modelling problem is limited to a maximum of three dimensions, corresponding to the 3D spatial dimensions $\mathbf{X} = [\mathbf{r}, \mathit{pitch}, \mathbf{z}]$ in a fixed measurement plane where total pressure and temperature measurements can be taken. While this is considered small in a data-

driven context, the flow inside an axial machine is expected to exhibit high complexity with strong gradients.

The approach to MFGP implementation significantly depends on the characteristics of the specific problem at hand and includes considerations of dimensionality and complexity. In practical terms, implementing MFGP involves addressing two key considerations. Firstly, important modelling choices within the GP framework must be made to ensure the method's adaptability to various real-world engineering scenarios. Secondly, the computational implementation of the algorithm itself needs attention to handle the substantial computational demands of GPs while ensuring numerical stability.

3.3.1. Kernel selection and influence

In line with the discussion in the previous chapter, particularly in Section 2.4.1, the primary modelling choice in GP implementation revolves around the selection of the kernel function. To guide this selection, the study conducted by KIANIFAR and CAMPEAN, 2020 serves as a valuable starting point.

In their comprehensive investigation, KIANIFAR and CAMPEAN, 2020 subjected 18 different test case functions to scrutiny, progressively increasing the complexity of these functions in terms of dimensionality. They also varied the dataset sizes and introduced artificial noise into the data. To evaluate various GP modelling kernel choices, they compared the GP models against alternative modelling approaches, such as polynomial fitting and Radial Basis Functions (RBF). Remarkably, regardless of the problem scale or the available dataset size, a GP model employing the Matérn-5/2 kernel consistently outperformed almost all other modelling methods, particularly when dealing with highly non-linear functions. Only in cases with high levels of artificially added noise did the Matérn-3/2 kernel exhibit greater robustness, showing performance comparable to the Matérn-5/2 kernel if the latter was slightly modified by introducing a bias kernel.

For the sake of robustness and wide-range applicability, the stationary anisotropic Matérn-5/2 kernel is selected. This choice is grounded in its demonstrated ability to capture complex patterns commonly encountered in engineering applications. The selection ensures that the resulting function is twice differentiable without imposing overly excessive smoothness, a characteristic sometimes observed in more default kernel choices but not often seen in real-world engineering scenarios.

The stationary nature of this kernel means that it depends solely on the distance between input locations $\mathbf{x}_i - \mathbf{x}_j$ and not on the specific values of \mathbf{x} . In simpler terms, it implies that the GP model's behaviour remains consistent throughout the input space,

regardless of where you are within that space. The stationarity of the Matérn-5/2 kernel ensures that the GP model remains flexible enough to adapt to different data patterns while maintaining the smoothness required for accurate modelling.

Additionally, the term "anisotropic" in this context means that the kernel is designed to account for varying degrees of correlation between input dimensions. In other words, it acknowledges that the relationships between input variables may not be the same in all directions. This is especially relevant in real-world engineering scenarios where certain input variables may exhibit stronger correlations or dependencies along one dimension than in others.

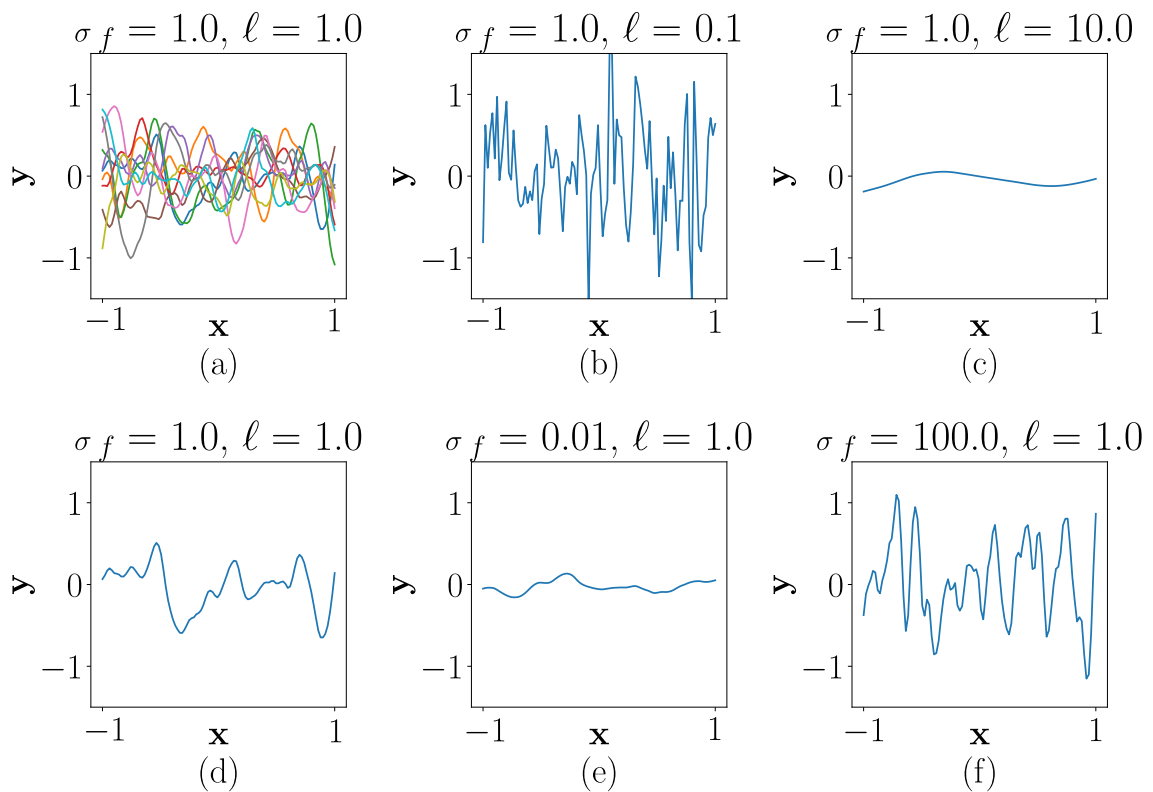


Figure 3.3. Prior samples for Matérn 5/2 kernel with different hyperparameter values, showcasing the effect of this parameters in GP modelling.

Consider the Matérn 5/2 kernel defined in equation 2.13, where ℓ represents the length scale for an individual input dimension. This parameter plays the role of determining the smoothness of the function being modelled. Essentially, it quantifies the extent to which a single data point affects neighbouring data points in the input space. Conversely, σ_f serves as a scaling factor, governing the vertical scale of the function relative to its mean.

To visually demonstrate the influence of these kernel hyperparameters on function modelling, Figure 3.3 is presented. This figure aims to elucidate the consequences of varying ℓ and σ_f while keeping the other parameter constant.

Figure 3.3 (a) showcases multiple samples from a reference kernel distribution with $\ell = 1$ and $\sigma_f = 1$. These multiple samples represent the inherent variability within the distribution, where a sample represents a complete function. Figures 3.3 (b) and (c) reveal how changing the length scale ℓ affects the function modelling while maintaining a constant scaling factor σ_f . Reducing ℓ to $\ell = 0.1$ results in a more oscillatory function due to the limited influence of each data point on nearby data points. Conversely, increasing ℓ to $\ell = 10$ leads to a smoother output as each data point's influence extends across a wider region.

In the bottom row of Figure 3.3 the effect of varying the scaling factor σ_f while keeping a constant length scale ℓ is investigated. Figure 3.3 (d) presents a single sample drawn from the reference distribution with $\ell = 1$ and $\sigma_f = 1$. In contrast, Figures 3.3 (e) and (f) display samples from distributions with $\sigma_f = 0.01$ and $\sigma_f = 100$, respectively. As previously discussed, σ_f controls the vertical variation of the GP model. Lower values of σ_f result in minimal deviations from the mean, while higher values, such as $\sigma_f = 100$, lead to pronounced smooth variations around the mean and typically indicate models with higher variability.

The selection of appropriate hyperparameters, such as the length scale ℓ and the scaling factor σ_f , play a significant impact in the behaviour of a GP model. These hyperparameters essentially define the model's behaviour and must be fit to match the characteristics of the available data and the specific problem at hand. The complete vector of hyperparameters in this work is:

$$\boldsymbol{\theta} = (\sigma_{f,high}, \ell_{1,high}, \dots, \ell_{N,high}, \rho, \sigma_{f,low}, \ell_{1,low}, \dots, \ell_{N,low}, \sigma_{y,high}) \quad (3.3)$$

3.3.2. Hyperparameters optimization

In the context of Gaussian Process modelling, the optimization of hyperparameters is typically carried out by maximizing the log marginal likelihood of equation 2.14. In practice, the negative of the log marginal likelihood is minimized instead in the optimization problem. This process essentially tunes the kernel hyperparameters to fit the observed data, resulting in a "trained" model that accurately represents the given data.

In this work, to find the optimal hyperparameters, the Limited-memory Broyden-

Fletcher-Goldfarb-Shanno with Bounds (L-BFGS-B) optimization method proposed by BYRD et al., 1995 is employed. The practical implementation of this algorithm is based on the open-source *Python* package *SciPy* developed by VIRTANEN et al., 2020.

The L-BFGS-B optimization algorithm is categorized as a quasi-Newton method, as it iteratively improves the optimization process by approximating the Hessian matrix. It does this by utilizing the analytical gradients of the log marginal likelihood equation with respect to the hyperparameters. The unique feature of L-BFGS algorithms is its low-rank approximation of the Hessian matrix based on a limited number of past iterations, which makes it highly memory-efficient.

L-BFGS-B is known for its fast convergence. However, the quality of the solution can be influenced by the choice of initial values. To ensure robustness, researchers often execute the optimization process multiple times with different initializations. In this work, the optimization process is restarted three times for each trained GP model.

L-BFGS-B is also designed to handle bounds on the optimization variables, ensuring that solutions remain within specified constraints. In the context of GPs, bounds are introduced to the hyperparameters before optimization. This seemingly subtle step serves a dual purpose: it enforces the discovery of a valid local minimum during optimization and acts as a prior on the hyperparameters, thus aligning with the Bayesian process.

In practical terms, lower bounds are typically applied to prevent hyperparameters from taking negative values. Additionally, an upper bound on the data noise is imposed to ensure that the log marginal likelihood does not attribute all data variability to noise. Moreover, in the context of a specific MFGP model, the scaling factor is bounded between 0 and 1, reflecting its modelling.

3.3.3. Computational implementation

The computational efficiency of the methodology heavily relies on the inversion of the term $(\mathbf{K} + \sigma_y^2 \mathbf{I})^{-1}$, which appears not only in the log marginal likelihood equation 2.14 but also in the predictive mean (equation 2.17) and covariance (equation 2.18). This matrix inversion operation has a computational complexity of $\mathcal{O}(N^3)$, where N represents the number of data points in the dataset. This computational cost can become prohibitively high for large datasets and is prone to numerical instability.

To tackle this challenge, the Cholesky decomposition method is employed. The Cholesky decomposition, also known as Cholesky factorization, is a technique that breaks down a positive-definite matrix into the product of a lower triangular matrix \mathbf{L}

and its conjugate transpose \mathbf{L}^T . In simpler terms, it simplifies a complex mathematical problem into more manageable steps. This decomposition is expressed as:

$$\mathbf{K} + \sigma_y^2 \mathbf{I} = \mathbf{L}\mathbf{L}^T \quad (3.4)$$

Once the Cholesky decomposition is performed, the matrix inversion problem can be simplified to:

$$(\mathbf{K} + \sigma_y^2 \mathbf{I})^{-1} = (\mathbf{L}^T)^{-1} \mathbf{L}^{-1} \quad (3.5)$$

The key advantage of this decomposition lies in its numerical stability and reduced computational complexity compared to other methods like direct matrix inversion. This enables efficient calculations of critical quantities in GP modelling, such as the log marginal likelihood, predictive mean, and predictive covariance. This efficiency is particularly relevant when dealing with large datasets, as is often the case in multi-fidelity modelling.

3.3.4. Implementation and example: Forrester function

In practice, the GP tools implemented in this work are adapted from open-source *Python* packages tailored for Gaussian Process modelling. The methodology relies on two main packages: *GPy* GPY, 2012 and *Emukit* PALEYES et al., 2019.

GPy, a Python-based Gaussian Process library developed by the Sheffield university machine learning group, provides essential functionalities for setting up GP kernels, integrate the tuning hyperparameters through optimization, and making predictions based on the trained models. *GPy* streamlines the complete GP modelling process. *Emukit* is used as "wrapper" of the functionalities of *GPy*, extending its functionality to accommodate the multi-fidelity framework.

These practical implementations empower the methodology with accessible and efficient computational tools and implementations that are needed to handle the complexities of real-world engineering scenarios, thus making the proposed approach a powerful solution for flow assessment in turbomachinery components while allowing its usage to any fluid mechanics area.

To illustrate the effectiveness of the MFGP and validate its implementation in this work, a small example is considered. The well-known 1D Forrester function, developed by FORRESTER et al., 2008 is assumed as the high-fidelity function:

$$f_{high}(x) = (6x - 2)^2 \sin(12x - 4) \quad (3.6)$$

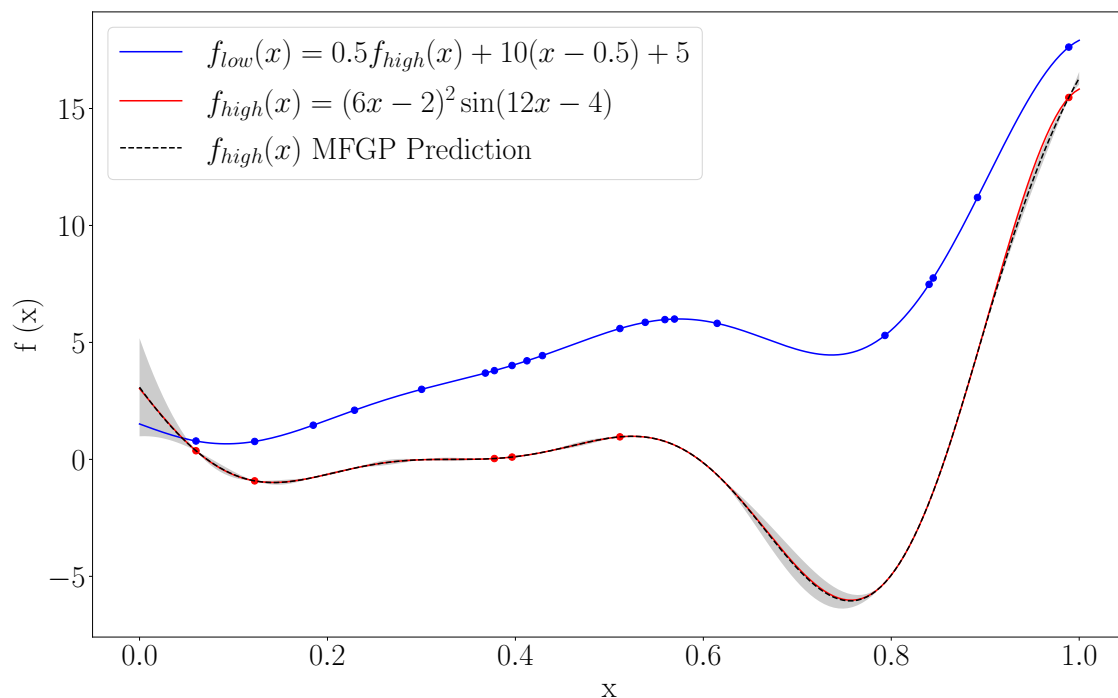


Figure 3.4. High- and low-fidelity Forrester functions and MFGP mean prediction and uncertainty of high-fidelity function. The dots represent the available observations. In this example, 20 low fidelity and 6 high fidelity samples are used in the MFGP model which predicts accurately the behaviour of the underlying truth.

For simplicity, it is assumed that sparse noiseless observations are available for this high-fidelity function.

A low-fidelity version of this function has also been proposed and here it is defined as:

$$f_{low}(x) = 0.5f_{high}(x) + 10(x - 0.5) + 5 \quad (3.7)$$

where the first term directly halves the high-fidelity function, the second term adds a linear shift and the third term adds a constant. Similar to the high-fidelity function, it is assumed that noiseless observations are available for the low-fidelity function, and these observations are available in greater quantity than those for the high-fidelity function.

Figure 3.4 illustrates both the underlying full low- and high-fidelity functions in blue and red, respectively. The dots represent the available observations. In this example, we use 20 low-fidelity and 6 high-fidelity samples to train the MFGP model.

The dashed black lines show the mean prediction and the associated uncertainty predicted by the MFGP implementation. Notably, the MFGP model captures the

behaviour of the underlying true functions remarkably well. Even in regions where high-fidelity data is not available, the model does an excellent job in predicting the function's trend.

One key area where the model performs correctly is the global function minima. Remarkably, the model accurately captures the location of the global function minima, even though it differs in location between the low-fidelity and high-fidelity versions of the function.

This example showcases the capabilities of MFGP in modelling complex functions using both high-fidelity and low-fidelity data. It demonstrates the potential of MFGP in accurately predicting functional behaviour and shows the robustness of the implemented approach.

3.4. Design of experiments: *A Priori* sampling

In the methodology schematic presented in Figure 3.1, two distinct approaches to DoE are introduced: the classic *a priori* sampling approach and an innovative *in situ* adaptive sampling approach rooted in the field of design of computer experiments and Bayesian Optimization, further explored in the next section.

In this section, the classic *a priori* DoE approach implemented in the methodology is discussed. Classic DoE methods predefine a fixed set of measurement locations prior to the actual data collection process. In essence, before any experimental data is acquired, both the number and positions of measurements are predetermined.

Selecting "where to measure" and "how much to measure" are fundamental decisions in any experimental campaign. Traditionally, these decisions have relied on a combination of expert knowledge, simple models, and preliminary tests. This is often the case because, at the planning stage, the actual experimental data is not available. Experimentalists also take into account practical considerations, including instrumentation limitations and access points to the flow under investigation.

Earlier in this work, in section 2.1, some approaches presented show preliminary tests for optimal probe placing, such as the work of SOUSA, GARCÍA-SÁNCHEZ, et al., 2018 while other showcased the influence of the probe placements in evaluating performance such as SESHADRI, A. DUNCAN, et al., 2022. This brings forth the question of not only about where and how much to measure but also about how to maximize the information extracted from these measurements.

With the demand of selecting not only where to measure but also how to position sensors optimally to extract the maximum information from the collected data, a num-

ber of sensor placement methods based on large amounts of data and with the goal of performing DA have been developed, usually by framing the problem in an optimization context as proposed by JOSHI and BOYD, 2009. Recently, leveraging big data trends, MANOHAR et al., 2018 propose a compressed sensing algorithm based on singular value decompositions that acts as a dimensionality reduction allowing for sparse sampling.

Recognizing the need to strategically position sensors to extract maximum information from the collected data, several sensor placement methods have emerged, often framed as optimization problems as in JOSHI and BOYD, 2009. Recently, MANOHAR et al., 2018 proposed a compressed sensing algorithm based on singular value decompositions, which acts as a dimensionality reduction technique, allowing for sparse sampling and efficient information extraction, leveraging large datasets to achieve better data assimilation.

Another literature approach has explored the integration of information theory and Bayesian approaches. Early work by SEBASTIANI and WYNN, 2000 employed Shannon entropy as a criterion for optimal experimental design. Building upon this foundation, KRAUSE et al., 2008 tackled the challenge of monitoring spatial phenomena using GPs by formulating an optimization problem to maximize the mutual information between measured and unmeasured locations. While highly informative, this approach involves an approximated combinatorial problem with considerable computational demands.

In more recent developments, TAJNAFOI et al., 2021 presented a data-driven approach that simplifies the complexity of big data problems. Their method begins by constructing a sparse GP model using extensive numerical datasets. This initial model informs subsequent sensor placement strategies, applying the algorithm introduced by KRAUSE et al., 2008 to strategically position sensors for optimal data collection.

The classic *a priori* DoE forms an fundamental building block of the proposed data-driven Bayesian hybrid measurement technique. The use of SVGP modelling is the integral part of the *a priori* DoE approach to precisely determine both the locations and the number of measurements required. This innovative approach couples SVGP with CFD RANS simulations to identify the optimal probe locations for flow assessment in real-world experiments.

3.4.1. SVGP as DoE tool: Implementation

As previously discussed in section 2.4.3, SVGP selects "inducing locations," denoted as \mathbf{Z} , to summarize a comprehensive dataset $y(\mathbf{X})$. This selection is accomplished by maximizing the ELBO of equation 2.27, or, in practical terms, minimizing it using a

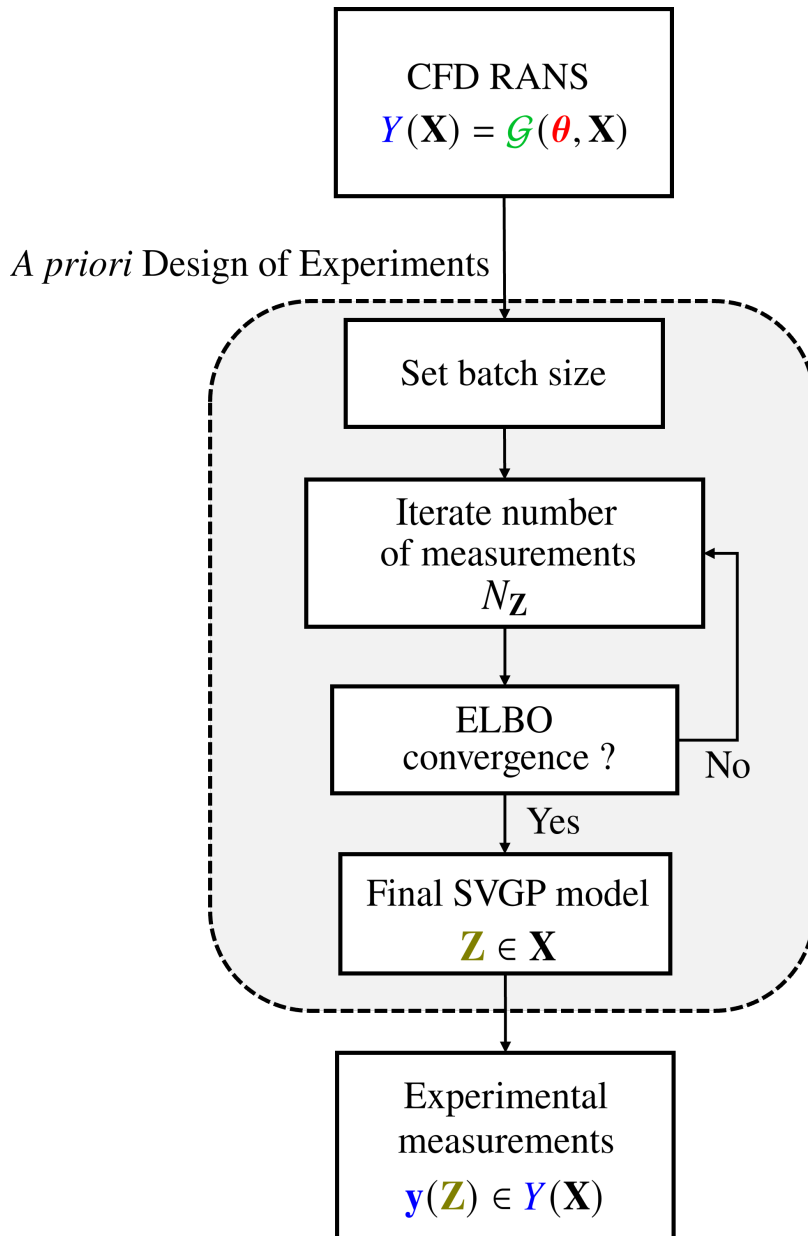


Figure 3.5. Detailed schematic of the proposed *a priori* DoE flowchart with SVGP modelling and CFD simulations to determine the optimal experimental measurement locations

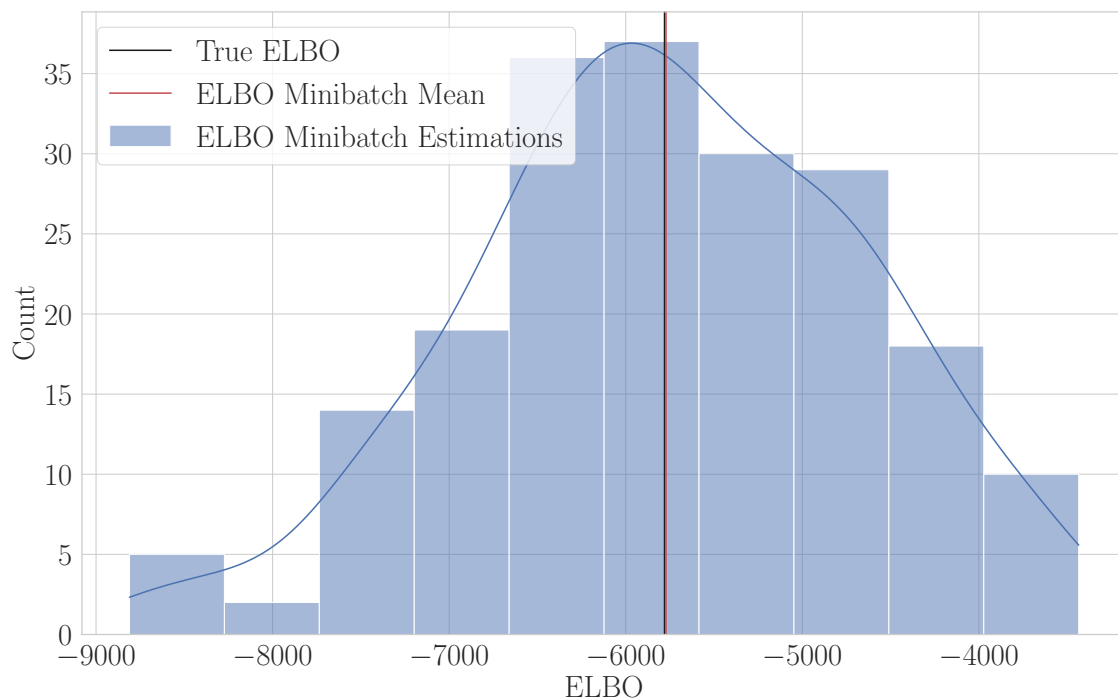


Figure 3.6. Histogram of ELBO evaluations using batch approach. The histogram mean is presented by a red line and the true ELBO value, computed with the full gradient is presented in black.

stochastic gradient descent optimizer. In this work, these "inducing locations" represent the proposed experimental measurement points, and the full dataset they summarize corresponds to the fully discretized CFD simulation.

The implementation of the *a priori* DoE approach with SVGP is based on the open-source Python package *GPy* GPY, 2012. The specific code used for SVGP was developed by the original authors of the SVGP modelling framework: HENSMAN et al., 2013, ensuring the fidelity and reliability of the implementation. This open-source foundation contributes to the transparency and accessibility of the methodology.

Figure 3.5 provides a detailed schematic of the DoE process with SVGP, illustrating the various steps essential for the robust application of the methodology.

To initiate the process, an available CFD RANS simulation is required. The first decision involves setting the batch size for optimization. This batch size determines the number of grid points from the CFD simulations used to compute the gradient in the negative ELBO optimization function. Typically, the batch size is determined through a trial-and-error process. The true value of the ELBO is computed at a fixed iteration, and the batch size is adjusted until convergence around the ELBO value is achieved.




Figure 3.7. ELBO optimization process for hyperparameter and "inducing locations" selection for the Forrester function. The ELBO evolution is shown for increasing number of "inducing locations". An ELBO upper bound cluster is obtained from $N_z = 100$.

To showcase the effectiveness of the batch size approach in estimating the ELBO accurately, Figure 3.6 presents a histogram of multiple evaluations of the ELBO using the batch approach. In this example, the Forrester low-fidelity function of equation 3.6 is used. The histogram's mean value is presented by a red line, while the true ELBO value, computed with the full gradient instead of mini-batches, is shown in black. The discrepancy between these lines illustrates the robustness of this approach in estimating the ELBO.

Subsequently, the need for selecting the number of measurements arises. In the literature, this decision is often based on computational cost considerations, given that SVGP models are utilized to manage extensive datasets, and each additional "inducing point" introduces complexity into the optimization process.

However, in this work, the computational cost argument is not deemed valid. It is assumed that the cost of extensive experimental measurements in an industrial context outweighs the computational expense.

An iterative procedure is employed to determine the optimal number of measure-

ments, denoted as N_z , necessary to capture the specific flow characteristics adequately. As the number of measurements increases, the ELBO also increases until an upper bound is found. At this stage, adding more "inducing locations" does not significantly enhance the ELBO.

To illustrate the iterative process done to select the optimal number of measurements based on the ELBO evaluation, once again the Forrester function is used as an example. Figure 3.7 illustrates the ELBO optimization process for hyperparameter and 'inducing locations' selection. It shows the ELBO evolution for increasing numbers of 'inducing locations,' and it can be seen at which point the ELBO reaches an upper bound.

From the lowest N_z value of 20 to 50, a noticeable decrease in the negative ELBO is obtained. This increase occurs again when N_z is increased to 100. However, when $N_z = 150$, the value around which the ELBO converges is approximately the same as when $N_z = 100$, as indicated by the overlap in the Figure. This suggests that in the SVGP model, increasing from 100 to 150 does not provide additional information to the model, and $N_z = 100$ is identified as the optimal number of 'inducing locations' for this toy example.

Another important point to mention is the non-smoothness of the ELBO convergence. This behaviour is justified by the discussion on batch size above. The fact that the optimization data changes from iteration to iteration leads to different gradient computations, but, as seen in Figure 3.6, on average, the ELBO converges to a correct value.

It is important to note that the ELBO convergence is not smooth. This behaviour is justified by the discussion on batch size above. The fact that the optimization data changes from iteration to iteration leads to different gradient computations, but, as seen in Figure 3.6, on average, the ELBO converges to a correct value.

This iterative process represents the most computationally intensive aspect of the DoE framework since multiple SVGP models must be trained. Additionally, it assumes that the flow patterns and their locations in the CFD simulations align with the real-world scenario, ensuring that future experimental measurements at selected points provide the same level of information as their corresponding locations in the CFD simulation.

Once the optimal number of "inducing locations" is determined, these locations \mathbf{Z} can be extracted from a final SVGP model. Subsequently, the experimental campaign can proceed to acquire measurements at these locations to provide data to the subsequent MFGP model and flow assessment.

3.5. Adaptive *In situ* sampling

While the classic *a priori* DoE approach proposed above is an excellent choice for industrial settings where cost efficiency and careful pre-planning of experimental campaigns are paramount, achieving precise flow assessment and the lowest level of uncertainty may necessitate additional experimental measurements in an academic environment.

In academic research, the pursuit of rigorous data and high precision justifies the allocation of more time and resources to experimental endeavours. In such scenarios, where the primary objective is to reduce uncertainty and gain deep insights into complex flow phenomena, the traditional *a priori* approach may prove to be less flexible and adaptive.

To meet the stringent requirements for achieving desired uncertainty levels within the proposed measurement technique, the concept of adaptive design of experiments, or *in situ* adaptive sampling is introduced. This approach is rooted in computational optimization studies, where new designs are iteratively selected based on both new and previous design to optimize the acquisition process. It leverages Bayesian Optimization (BO), which extends the Bayesian formalism to perform robust optimization.

The concept of BO is not novel in the realm of computer experiments and can be traced back to its initial introduction by SACKS et al., 1989. Over time, a substantial body of literature has evolved, addressing the design of computer experiments. A comprehensive review of this field is presented by GARUD et al., 2017.

Within this extensive literature, two distinct approaches have emerged: one-shot DoE and adaptive sampling. The one-shot DoE typically involves the use of 'static' designs, often referred to as space-filling designs. A prime example of this is Latin Hypercube Sampling (LHS) designs, pioneered by JIN et al., 2005. These designs align closely with the *a priori* sampling discussion presented earlier in this section.

In contrast, the adaptive sampling approach, closely related to the *in situ* sampling required for this work, takes a dynamic stance. It continually selects measurement locations and quantities as the experiment unfolds, capitalizing on the latest experimental data to enhance flow assessment.

The selection of the next measurement location is guided by a careful balance between estimated error and model uncertainty. This strategy enables to strike an optimal equilibrium between the desire for high-quality data and the constraints of time and resources. To the authors knowledge, this represents the first instance of adaptive sampling being used in a full experimental context.

The cornerstone of *in situ* adaptive sampling lies in the selection of the next mea-

surement location. This process demands a precise balance between two fundamental strategies: global exploration and local exploitation. Achieving this equilibrium is paramount for accurate flow assessment.

Global exploration entails seeking measurement locations in regions of the experimental domain that have not yet been adequately sampled. This strategy is essential for capturing the broader characteristics of the flow field, especially in flow regions that may be undersampled in the initial measurements. By actively exploring these unexplored regions, the adaptive approach can reveal hidden insights and potential improvements in model accuracy.

On the other hand, local exploitation targets specific areas of interest, aiming to refine the model's understanding of complex flow phenomena and reduce uncertainty in those regions. It implies taking measurements in regions where the model already demonstrates a degree of certainty but can benefit from finer granularity. Exploitation aims to enhance the model's accuracy in areas where it's already reasonably well-informed but strong gradients are present.

Implementing an appropriate algorithm to weight these two aspects is the key step in adaptive sampling. If global exploration is prioritized too much, the final flow model might present high uncertainty in high gradient regions since it should have been refined more. Conversely, focusing too heavily on local exploitation may hinder the sampling of certain flow regions.

The review by FUHG et al., 2021 provides a comprehensive comparison of various adaptive sampling algorithms within the context of GPs. The study breaks down these algorithms according to their strategies for both exploitation and exploration, conducting evaluations across 16 different test functions of varying complexity, taking into consideration factors such as the optimization goal, functional response expectations, data availability, dimensionality, and computational cost.

Following a meticulous assessment, the adaptive sampling approach selected for determining the next measurement point in the *in situ* DoE is the Maximize Expected Predicted Error (MEPE) strategy. This methodology, introduced by H. LIU et al., 2017, relies on quantifying the Expected Predicted Error (EPE) of a GP model. The EPE comprises two components: a bias term and a variance term, which correspond to local exploitation and global exploration, respectively.

The bias term is concerned with regions where significant prediction errors exist, thereby emphasizing local exploitation. Unfortunately, it cannot be directly computed due to the unavailability of prior data. To address this, the authors approximate the bias term using a leave-one-out cross-validation (LOOCV) error, employing the effi-

cient computational method established by SUNDARARAJAN and KEERTHI, 2001. To ensure continuity, this approximation assumes that the LOOCV error at an unobserved location \mathbf{x} is equivalent to the error at its nearest available sample, denoted as $\hat{e}_{\text{LOOCV}}^{\text{approx}}$.

However, this approximation of the bias term tends to introduce bias in practical cases. To mitigate this, a balance factor α is introduced, enabling a trade-off between local exploitation and global exploration. The EPE criterion is defined as a weighted combination of these terms:

$$\text{EPE}(\mathbf{x}) = \alpha \hat{e}_{\text{LOOCV}}^{\text{approx}^2}(\mathbf{x}) + (1 - \alpha) \sigma_{\mathbf{y}}^2(\mathbf{x}) \quad (3.8)$$

The balance factor α adapts based on the previous iteration's error approximation:

$$\alpha = \begin{cases} 0.5, & \text{if } m_0 = 0 \\ 0.99 \min \left[0.5 \frac{e_{\text{true}}^2(\mathbf{x}_m)}{e_{\text{LOOCV}}^2(\mathbf{x}_m)}, 1 \right], & \text{else} \end{cases} \quad (3.9)$$

Here, m signifies the iteration number, and in the $m + 1$ iteration, the newly acquired measurement at sampling location x_m is used to compare the true GP model error of the previous iteration, denoted as $e_{\text{true}}^2(\mathbf{x}_m)$, with its original estimation, $e_{\text{LOOCV}}^2(\mathbf{x}_m)$. This comparison defines the balance between exploration and exploitation.

Finally, the next sample point is determined by maximizing the EPE across the input space:

$$\mathbf{x}_{\text{MEPE}}^{m+1} = \arg \max_{\mathbf{x} \in \mathbf{X}} (\text{EPE}(\mathbf{x})) \quad (3.10)$$

This approach encapsulates the essence of intelligent measurement selection within the *in situ* DoE, ensuring a balanced pursuit of global exploration and local exploitation to attain precise flow assessment.

Having defined MEPE as the adaptive sampling scheme for the in situ sampling approach proposed in the methodology of this work, Figure 3.8 illustrates a detailed schematic of the flowchart of the in situ sampling approach.

The process starts with an initial set of experimental measurements $y(\mathbf{x}_{m_0})$ taken a priori in sampling locations \mathbf{x}_{m_0} . This initial set of measurements is typically small, with some authors suggesting a size of 10D, where D is the dimensionality. This choice is based on practical considerations, aiming to balance the need for an informative initial dataset. Alternatively, this initial set can be determined using the SVGP approach proposed in the previous section or with a space-filling LHS design for simplicity.

The adaptive sampling approach iterative process then commences, continuing until predefined stopping criteria are met.

A GP model is trained with the available experimental measurements. A standard GP model is employed here instead of a MFGP model. The choice to use a standard GP model is driven by the substantial computational overhead associated with fitting an MFGP model at each iteration.

The stopping criteria are evaluated based on the GP model's predictive mean and covariance. Typically, these criteria are set based on the desired level of accuracy or uncertainty reduction, ensuring that the adaptive sampling process continues until the specified information goals are met.

At each iteration, until the stopping criteria are satisfied, the MEPE adaptive scheme, as previously described, is applied. The Expected Predicted Error (EPE) acquisition function is computed for all candidate locations \mathbf{x} within a predefined domain space \mathbf{X} . The candidate measurement location where the EPE is maximized becomes the next measurement location \mathbf{x}_{m+1} to acquire.

Once a new measurement $\mathbf{y}(\mathbf{x}_{m+1})$ is experimentally acquired, this new information is incorporated into the available data. Subsequently, a new GP model is fitted to the expanded dataset, and the process continues iteratively until the stopping criteria is achieved. This completes the *in situ* sampling DoE and the proposed methodology process continues with the training of the MFGP described earlier.

Incorporating the *in situ* adaptive sampling approach, driven by the Maximize Expected Predicted Error (MEPE) criterion, as an option to the proposed methodology offers an alternative to a precise flow assessment in complex academic research settings. This approach, emphasizing a balance between global exploration and local exploitation, allows to continually enhance flow assessment by leveraging newly acquired experimental data. Through the iterative selection of measurement locations, considering model error and uncertainty, the collection of high-quality data while effectively managing time and resource constraints is ensured. This application of advanced sampling strategies within a full experimental context is a novelty introduced in the proposed methodology to exploit at maximum the need for less measurements while accurately assessing the flow field.

3.6. Methodology Performance Evaluation

The performance of any methodology hinges on its ability to provide accurate and dependable results. In the context of the proposed flow assessment methodology, a rigorous performance evaluation is essential. This evaluation will involve applying the methodology to two test cases while using separate validation flow field datasets that

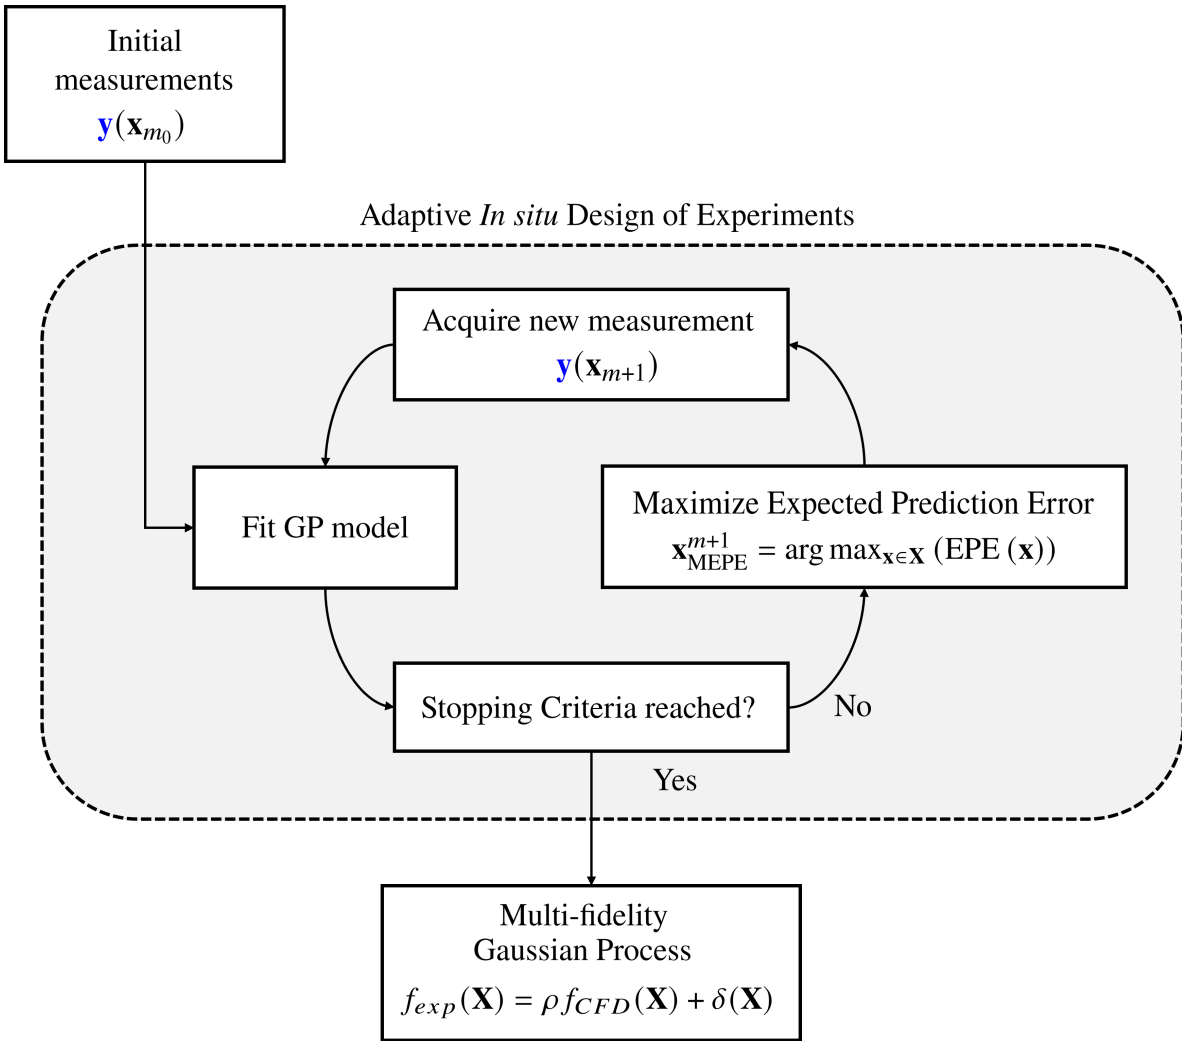


Figure 3.8. Detailed schematic of the proposed adaptive *in situ* DoE with the MEPE sampling approach selecting iteratively where to acquire the next experimental measurement.

will not be used in the methodology's flow assessment.

This section introduces the key parameters and metrics that will comprehensively evaluate the methodology's performance, serving as critical indicators of its effectiveness and reliability.

To assess the methodology's precision in predicting flow characteristics, three fundamental metrics will be evaluated:

Root Mean Squared Error (RMSE): RMSE is a standard error metric that quantifies prediction accuracy. It calculates the square root of the average squared differences between predicted and actual values.

$$\text{RMSE} = \sqrt{\frac{1}{N_{val}} \sum_{i=1}^{N_{val}} (\mathbf{y}_*^i - \mathbf{y}_{val}^i)^2} \quad (3.11)$$

where \mathbf{y}_*^i represents the predicted values and \mathbf{y}_{val}^i represents the actual validation data.

Maximum Absolute Error (MaxAE): MaxAE identifies the most significant prediction error, allowing a direct comparison with experimental measurement uncertainty. A notable difference between MaxAE and RMSE indicates problematic model performance in specific flow regions.

$$\text{MaxAE} = \max_{1 \leq i \leq N_{val}} |\mathbf{y}_*^i - \mathbf{y}_{val}^i| \quad (3.12)$$

R-squared (R^2): R-squared evaluates the proportion of the variance in the dependent variable (flow field) that can be explained by the independent variables (model predictions). A higher R^2 value indicates a better fit of the model to the data.

$$R^2 = 1 - \frac{\sum_{i=1}^{N_{val}} (\mathbf{y}_*^i - \mathbf{y}_{val}^i)^2}{\sum_{i=1}^{N_{val}} (\mathbf{y}_{val}^i - \bar{\mathbf{y}}_{val})^2} \quad (3.13)$$

where the $\bar{\mathbf{y}}_{val}$ is the mean of the actual validation data

In an experimental setting, accurate flow assessment requires proper uncertainty quantification. This methodology inherently incorporates Bayesian formalism, offering an evaluated uncertainty with the flow predictions. The MFGP approach directly provides an uncertainty evaluation by considering the square of the diagonal elements in the predictive covariance matrix.

In the following chapters, the two test cases used to showcase the proposed methodology's performance are presented.

Chapter 4.

Test Case - H25 Axial Compressor

THE application of the methodology to assess the flow of the H25 axial compressor is presented in this chapter. This test case is taken as a reference test case. The experimental and numerical tools used to support the application of the methodology are initially described. The MFGP modelling approach is showcased with random experimental sampling to then evaluate the impact of the proposed DoE approach. Finally, the *in situ* sampling approach is assessed.

Overview

4.1. Introduction	61
4.2. Experimental tools	62
4.3. Numerical tools	65
4.4. Results	67
4.4.1. Preliminaries	70
4.4.2. Multi-fidelity Gaussian Process	70
Pressure Field Mean Flow Assessment	71
Sensitivity Analysis	74
Pressure Field Uncertainty Analysis	76
Temperature Field Assessment	78
Conclusions	79
4.4.3. <i>A priori</i> Design of Experiments	79
Selecting Number of Measurements	80
SVGP Sensitivity Analysis	80
Conclusion	84

4.4.4. *In situ* Sampling 84
MEPE Acquisition Function Behaviour 85
Adaptive Sampling Sensitivity Analysis 91
Initial Dataset Influence 93
Stopping Criteria 95
Conclusion 97



Figure 4.1. H25 compressor blades

Table 4.1. H25 stage design parameters.

Parameter	Value
Blade height	25 mm
Hub-tip ratio	0.9
Rotor blades	65
Stator blades	80
Pressure ratio	1.25
Nominal speed	9650 rpm

4.1. Introduction

Selecting a suitable reference test case is a crucial decision in assessing the robustness and effectiveness of a methodology. In such a choice, there is a delicate balance to be struck between real-world representativeness and the availability of comprehensive datasets for sensitivity analysis.

A reference test case should represent the industrial context, particularly one susceptible to instrumentation issues. Simultaneously, it should allow a full exploration of the methodology's implementation details, making sensitivity analysis feasible. Such exploration often demands access to extensive datasets that, while simpler than real-world experimental flows within turbomachine components, facilitate rigorous testing of the methodology.

The H25 axial compressor is a research stage developed in the frame of the LEM-COTEC project (FP7 - Low Emission Core Engine Technologies) to evaluate core size effects on high-speed axial compressor performance and flow field. It is representative of the last stage of a multi-stage axial compressor for aircraft gas turbine applications and features a low aspect ratio (span-to-chord ratio) and a highly three-dimensional blade design. It features an exceptionally reduced blade height, from which it derives its name, with the blade height measuring a mere 25mm. Figure 4.1 shows the H25 compressor stage with its 25 mm height blades. For the sake of completeness, Table 4.1 presents the design parameters of the H25 compressor.

The configuration of the H25 axial compressor, characterized by its reduced blade height, makes it an ideal candidate for the methodology's assessment since it aligns with the research goals of improving experimental flow assessment methods in next generation compressor technologies that are more sensitive to instrumentation since their size is decreasing and probe minimization cannot follow.

More importantly, data availability played a major role in the selection of this test

case. The H25 axial compressor has been extensively and comprehensively experimentally tested by BABIN, 2022, providing access and in-depth knowledge of the flow. This data includes a wide array of measurements related to the flow characteristics of the compressor that are described below.

The availability of such a comprehensive dataset allows for rigorous testing and validation of the methodology's different components. By employing this test case, it is possible to independently analyse and validate each step of the methodology, ensuring its reliability and effectiveness under real measurement conditions.

Furthermore, the H25 axial compressor data availability makes it an ideal reference test case for evaluating the sensitivity of different factors, such as the quantity and location of experimental measurements

These factors collectively make the H25 axial compressor an ideal test case for the methodology presented in this work. In essence, the selection of the H25 axial compressor test case is driven by the unique opportunity it presents for rigorous testing, validation, and sensitivity analysis while it also exemplifies a real fluid mechanics test case, emphasizing the methodology's applicability in contemporary engineering problems.

In the following sections, the experimental and numerical tools used to support the application of the methodology to this test case, and subsequently the results of it are presented.

4.2. Experimental tools

The experimental data was acquired in the R4 test rig of the von Karman Institute for Fluid Dynamics (VKI). The layout of this facility is depicted in Figure 4.2. This facility incorporates an inlet plenum equipped with a heat exchanger and honeycomb (1), which regulate the total temperature and flow homogeneity. The flow is then channelled to the H25 test section (2), via a smooth convergent bell-mouth. After passing through the test section, the compressed flow is released into a large collector (3). A return duct (4) connects this collector back to the inlet plenum, forming a closed-loop system. The facility includes a precision throttle valve (5) that enables accurate control of mass flow.

The R4 facility operates in a closed-loop configuration, providing the flexibility to independently adjust key parameters, including Reynolds number (Re), rotational speed, and throttling transients. All experimental data used for this work was obtained under atmospheric conditions.




Figure 4.2. VKI R4 facility layout for axial compressor studies.

A meridional view of the H25 test section, labelled as point (2) in Figure 4.2, is presented in Figure 4.3. In this view, the rotating components are highlighted in red, while the stator row is marked in green. The instrumentation is axially distributed along four measurement planes, which are referenced as MP (Measurement Plane). MP0: Located after the convergent inlet bell-mouth, seven rotor chords upstream of the rotor. MP1: One chord upstream of the rotor. MP2: In the inter-row region between rotor and stator. MP4: At the stage outlet, one chord downstream of the stator blades.

The overall performance, particularly pressure ratio and efficiency, is assessed using combined total pressure-total temperature rakes deployed at MP0 and MP4. Additionally, radial traversing of probes is possible in all measurement planes.

At MP4, a motor was installed to allow a pitch-wise traversing of the probes over one and a half stator pitch. This setup enabled the sampling of a radial-azimuthal two-dimensional map of the flow field. The acquisition process for this outlet flow cartography, while essential for evaluating the flow inside the compressor, was the most time-consuming part of the experimental campaign, taking approximately 3 hours to collect data points from 950 locations. For this reason, and given the critical role of MP4 in assessing the compressor's performance, this experimental test at the machine's design point serves as the reference case for applying the proposed methodology. Figure 4.4 illustrates the experimental grid employed in this test, with 950 grid points distributed in a matrix of 25 radial points over 38 pitch-wise points, featuring refine-




Figure 4.3. H25 compressor stage test section meridional view

ment near the walls in the radial direction for improved gradient evaluation and uniform discretization in the tangential direction.

For this test, a custom-designed miniaturized three-hole probe was utilized to measure flow total pressure, angle, and Mach number. The systematic uncertainty of the miniaturized pressure probe is determined by the pressure scanner used. In this case, it amounts to $\pm 0.08\%$ of the full span. For a measurement in MP4 this corresponds to ± 84 Pa. In terms of total pressure ratio measurements relative to inlet conditions (MP0), where the systematic uncertainty is estimated to be ± 14 Pa, the propagated systematic total pressure ratio uncertainty, computed with a truncated Taylor series expansion, is ± 0.00088 . Based on the time series variation of the measured pressure ratio, at a constant throttle operation, the random error associated with the pressure ratio measurement is estimated. Thus, the overall pressure ratio is ± 0.00133 .

A similar miniaturized thermocouple probe was designed in-house for measuring flow total temperature. A closer view of this probe is presented in Figure 4.5a, and its installation in the compressor stage, highlighting the reduced scale of the instrumentation and the compressor stage itself, is shown in Figure 4.5b. The estimated propagated uncertainty for total temperature measurements is ± 0.6 K.

The probes were installed with a standard mechanical safety margin of 0.02 mm to ensure a gap between the probe and hub/casing walls. Additionally, the radial size of the probe head was added to this value when determining the accessible span for each probe. This led to a 2% span end-wall proximity for the pressure probe, with the last





Figure 4.4. H25 experimental test grid at MP4.

measured temperature point located at 4% span from the nearest end wall due to its larger probe head. For more details, refer to BABIN, 2022.

4.3. Numerical tools

With a comprehensive understanding of the experimental setup and data acquisition, another essential tool required for applying the methodology proposed in this work, the CFD numerical model of the test case is introduced.

The primary goal of the numerical simulations is to harness the highly discretized flow



(a) Probe zoom. (b) Probe installed in compressor stage. Blades blurred.

Figure 4.5. Miniaturized total temperature probe.




Figure 4.6. H25 compressor stage numerical CFD domain.

information obtained from them in conjunction with the undersampled experiments, as outlined in the proposed hybrid methodology. To achieve this, it is critical to ensure that the simulation domain aligns with the experimental setup.

This domain extends from experimental Measurement Plane 0 (MP0) to a location two rotor chords downstream of MP4. This extension helps avoid potential effects related to outlet boundary conditions at the experimental stage outlet (MP4). Coupling numerical and experimental data at this location is fundamental as it serves as the reference test case for assessing the methodology’s performance.

Technological features, such as blade fillets, have been retained in both the rotor and stator sections of the numerical domain. However, cavities present in the experimental setup have been excluded. Under the assumption of periodicity, only one blade passage is simulated. A mixing plane approach is employed to connect the rotating rotor block with the non-rotating stator block, ensuring a conservative coupling of pitchwise quantities.

For mesh generation, a structured grid was created using Autogrid 5. Figure 4.6 presents both a meridional view and a blade-to-blade zoom of the computational grid. An O-4H mesh topology was adopted for both blade rows, while an OH topology was

used in the rotor tip gap region. Notably, all values of the non-dimensional wall distance (y^+) were maintained below 2.

The mesh generation involved a multi-block structured grid approach, resulting in three mesh levels. Grid convergence analysis was conducted across these three levels, which comprised grids with 0.6 million cells, 4.55 million cells, and 19.45 million cells, following the methodology of CELIK et al., 2008. The Reynolds-Averaged Navier-Stokes (RANS) data used in this work were obtained using the mesh with 4.55 million cells.

NUMECA FINE Turbo 12.1 was used to solve the fully turbulent compressible steady RANS equations. The k - ω SST turbulence model MENTER, 1994 was selected as the turbulence closure model.

As previously mentioned, the chosen operating regime for this study was the design point of the compressor, specifically at nominal speed. At the inlet of the numerical domain (corresponding to MP0), the experimentally measured pitchwise averaged total pressure and total temperature profiles were imposed. At the numerical domain's outlet, the absolute mass flow is imposed. Additionally, an adiabatic assumption is made for all the walls within the domain.

4.4. Results

In this initial result section, a comprehensive comparison between the experimental flow data collected in the H25 stage outlet and the corresponding numerical results obtained using the CFD model is presented, specifically in terms of total pressure and total temperature. The data collected at MP4 holds a central position in this work, serving as the reference test case for the application of the proposed methodology.

The primary objective of this comparative analysis is to evaluate and emphasize the agreement or potential discrepancies between different fidelity data sources—experimental and numerical. Since the H25 compressor is the reference test case in this work, it is important to understand how the level of correlation between these data sources influences the performance of the proposed methodology.

Therefore, this comparison serves as a fundamental basis for the subsequent analysis of the methodology. It ensures that the methodology can effectively harness the numerical model to enhance flow assessment, especially in scenarios where experimental data might be limited or otherwise unavailable.

Figure 4.7 provides a visual comparison of the normalized total pressure flow field between the high-fidelity reference experimental dataset (on the left) and the low-fidelity CFD RANS numerical flow (on the right), limited radially by the probe sampling con-




Figure 4.7. H25 stage outlet Total Pressure map comparison between experimental measurements (left) and CFD simulation (right).

straints.

The comparison reveals that the overall flow topology remains consistent between the two datasets, with secondary flow structures clearly visible in both. However, a clear discrepancy emerges in terms of pressure values, resulting in an pointwise average pressure ratio difference of around 6% between the experimental and numerical datasets.

Examination of the flow fields indicates that the RANS simulation tends to over-predict the depth and radial thickening of the wake. This over-prediction can be attributed to the radial streamline contraction caused by the highly three-dimensional blade design of the stator. The stator blades exhibit an high lean angle at the hub region, leading to a strong transverse pressure gradient over the blade (as demonstrated in TAYLOR and MILLER, 2017).

The application of a mixing plane approach at the rotor-stator interface results in the tangential averaging of the flow, thereby altering the stator’s inlet conditions, especially in terms of blade incidence. This primarily affects the regions close to the hub and tip, where the CFD simulation tends to over-predict pressure losses and the size of secondary flow structures. This discrepancy is most evident at the casing, where secondary wall flows merge with the wake, and at the hub, where the signature of the passage vortex is clear, while in the experimental data, it is believed that the passage vortex’s size is likely beyond the probe sampling constraints.

A comparison analysis is conducted for the total temperature 2D map, as presented in Figure 4.8. This figure illustrates a comparison between the high-fidelity reference experimental total temperature map (on the left) and the low-fidelity CFD RANS numerical total temperature flow (on the right). The numerical flow data is once more




Figure 4.8. H25 stage outlet normalized Total Temperature map comparison between experimental measurements (left) and CFD simulation (right).

constrained within the radial boundaries determined by the probe sampling.

In average absolute terms, a difference of 0.5% is observed between the two datasets. However, visually, it becomes evident that the flows exhibit a high mismatch. The discrepancy in the total temperature flow mainly arises from the adiabatic wall assumption made in the numerical model.

Specifically, the experimental dataset displays higher temperatures in the hub region, directly downstream of the blade, particularly around $\theta=1$. This elevated temperature is attributed to heat transfer effects from the rotating machinery, including the rotating shaft and bearings, to the flow. These thermal interactions are not present at all in the numerical model, resulting in lower temperature values in this area.

A similar heat transfer analysis can be extended to the casing region. In the experimental setup, the compressor dissipates heat into the colder environment through the tip casing. However, this heat dissipation is not accurately modelled in the numerical simulation. Consequently, the numerical model indicates an accumulation of heated fluid at higher radial regions, leading to temperature differences between the two datasets.

These discrepancies are not indicative of a failure in the methodology or the experimental procedures but rather arise from the fundamental modelling choices in RANS simulations, which provide the low-fidelity data that complements the under-sampled high-fidelity experimental data. These observed differences are a direct result of the modelling choices that underpin RANS simulations.

It is essential to recognize that RANS serves as the primary source of low-fidelity data in our methodology. These low-fidelity data sources complement the high-fidelity

experimental data to enhance the flow assessment process, particularly in cases where experimental data may be limited or under-sampled. Understanding these modelling nuances is the initial step in applying the proposed methodology effectively.

In the next section, the proposed methodology data fusion algorithm, multi-fidelity Gaussian Process is applied to the above reference flow fields to showcase the applicability of this algorithm in the proposed methodology for flow and uncertainty assessment.

4.4.1. Preliminaries

The following section focus on the proposed methodology capability to merge information from two distinct fidelity data sources: high-discretized CFD RANS low-fidelity data ($\mathcal{D}_{low} = (\mathbf{x}_{low}, \mathbf{y}_{low})$) and under-sampled high-fidelity experimental data ($\mathcal{D}_{high} = (\mathbf{x}_{high}, \mathbf{y}_{high})$), where $\mathcal{D}_{low} \gg \mathcal{D}_{high}$.

Recalling the notation used in Chapter 2, the outlet flow assessment involves taking the physical domain as an input location, $\mathbf{x} = [\mathbf{r}, \boldsymbol{\theta}]$, which is mapped to either the outlet total pressure, $\mathbf{y} = \mathbf{P}_{T4}$, or the outlet total temperature, $\mathbf{y} = \mathbf{T}_{T4}$, for this specific test case. Throughout the following analysis, the input domain \mathbf{x} is normalized between 0 and 1. Additionally, flow quantities are normalized with respect to a reference value.

Special care is taken to ensure the high and low fidelity datasets match in the input domain. This requires to shift the available datasets since the CFD RANS simulation domain '0' does not match the pitch wise '0' of the experimental test. In the simulation domain, the blade is in the centre of the studied passage while for the experimental test, the probe is initially aligned with the mechanical zero of the blade. Based on this information, domain matching is possible. This ensures that the MFGP model infers the correct relation between fidelities. Moreover, the domain matching is required to obtain a correct grid from the proposed DoE based on SVGP.

The complete available dataset is divided into two distinct subsets: the *training* subset and the *validation* subset. The *training* subset is utilized in the optimization process, maximizing the marginal log-likelihood of Equation 2.14, enabling hyper-parameter learning and flow reconstruction. The *validation* subset, serving as the ground truth, is afterwards compared against the methodology predicted flow reconstruction to validate and evaluate its performance.

4.4.2. Multi-fidelity Gaussian Process

To evaluate the proposed methodology, an initial step is to independently evaluate the MFGP framework as a data fusion algorithm. In this section, the data is randomly split

into these subsets, with various percentages of subsampling. Specifically, subsampling percentages of 10%, 20%, 33%, and 50%, corresponding to 95, 190, 313, and 475 absolute measurement points, respectively are explored.

Pressure Field Mean Flow Assessment


Figure 4.9 presents a comparison between the complete experimental reference data (on the left) and the flow reconstruction achieved using the MFGP model (on the right) for three different measurement percentages: 10% (4.9a), 20% (4.9b), and 33% (4.9c). The white crosses on the experimental reference represent the randomly selected $\mathcal{D}_{low,train}$ experimental measurements used as high-fidelity samples for multi-fidelity model training. The scaling factor (ρ), that relates the RANS simulation to the experimental data, is also provided.

An initial qualitative analysis of the flow fields reveals a generally consistent match in critical flow features and their locations, with an increasingly improved visual alignment as the number of measurements increases.


When only 10% of the measurements are used, the secondary flow structures show discrepancies in both shape and magnitude compared to the full dataset. However, their locations align closely. This is particularly evident in the wake region and the casing secondary structure. In these areas, where high-fidelity measurements are scarce due to random undersampling, the MFGP flow assessment tends to favour the highly discretized, low-fidelity RANS flow presented in Figure 4.7 since not enough high-fidelity information is available to describe these regions.

With an increased use of experimental data, specifically 20% of the complete dataset, the MFGP model provides a 25% more accurate prediction of the entire experimental flow than with 10% of the measurements. It correctly assesses high-pressure regions and secondary flow structures in terms of both shape and location. Notably, the model accurately captures the radial contouring of the strong gradient wake, a detail not accurately matched in the CFD simulation. Additionally, other secondary flows, such as the passage vortex near the boundary layer hub region and losses emanating from the rotor tip vortex in the casing, are also visibly and correctly predicted in terms of size and magnitude.


To accurately capture the depth of the wake, a higher percentage of measurements (33%) is required. It is at this percentage of data that the MFGP correctly assesses the entire flow, including the low-pressure pocket within the wake. Moreover, the MFGP model tends to smooth the flow since it is not limited to the probe spatial discretization and it can be evaluated in a highly discretized manner, opening the



(a) 10% Measurements used. $\rho = 0.374$.



(b) 20% Measurements used. $\rho = 0.158$.



(c) 33% Measurements used. $\rho = 0.110$.

Figure 4.9. MFGP total pressure flow assessment at MP4 (right) with different experimental measurements undersampling percentages against fully sampled experimental reference (left).

Table 4.2.: Experimental hyperparameters and performance metrics for the different models presented in Figure 4.9.

Measurements	ρ	σ_f	ℓ_{pitch}	ℓ_r	σ_y^2	RMSE	MaxAE	R^2
10%	0.374	0.214	2.503	3.721	4.8E-06	3.1E-03	1.7E-02	0.900
20%	0.158	0.288	1.433	1.252	9.4E-08	1.2E-03	1.2E-02	0.984
33%	0.110	0.309	1.540	1.276	5.1E-08	4.5E-04	1.6E-03	0.998

possibility of obtaining super-resolution.

In summary, the MFGP’s ability to align high- and low-fidelity data seems to improve as more high-fidelity measurements are included. With a higher percentage of measurements, but still with an high degree of undersampling (33%), it provides a comprehensive assessment of the flow, closely matching the reference experimental results from a visual point of view.

To further support this analysis, a quantitative evaluation of the performance metrics introduced earlier in section 3.6 is presented. These metrics are summarized in Table 4.2 alongside the estimated hyperparameters for the specific models presented in Figure 4.9.

A quantitative analysis from the table offers similar insights as the visual analysis described above. The scaling factor (ρ), that relates the RANS simulation to the experimental data, decreases as the percentage of experimental measurements increases, being nearly four times higher for 10% compared to 33%, meaning that the relevance of the RANS simulations is decreasing. This justifies the clear flow signatures from the RANS simulation dominating in regions where no experimental data is available with only 10% measurements.

A convergence of hyperparameters is observed, with values being nearly equal for the 20% and 33% cases. These slight differences correspond to the assessment of the core of the wake. The MFGP model’s estimated experimental uncertainty (σ_y^2) decreases with an increase in measurements, this suggests that, as data increases the MFGP model does not require to explain data with noise.

The error metrics, such as RMSE and MaxAE, show a consistent trend: they decrease by an order of magnitude from 10% to 33% measurements. These values approach the estimated experimental uncertainty of $\pm 1.3E-03$, indicating an improved representation of the true experimental flow.

On one hand, when high-fidelity experimental measurements are available to the model, the error and uncertainty are lower, on the other hand, the goal of this work remains to propose an hybrid experimental methodology that provides an accurate flow




Figure 4.10. RMSE of Total pressure ratio of across 50 MFGP models. Sensitivity to the random selection of experimental data and the percentage of experimental data used in the model training.

prediction with the maximum reduction in measurements required and testing time, meaning, reducing the number of high-fidelity samples to a minimum. Intuitively, a trade-off relation can be found between the number/location of the experimental high-fidelity measurements used and accuracy in the mean flow prediction and standard deviation.

Sensitivity Analysis

To further investigate this trade-off, the availability of a full reference test is exploited to perform a sensitivity analysis. This analysis assesses the impact of two crucial factors on the performance of the proposed methodology using error metrics. The effect of randomly selecting the *training* measurements subset from the full dataset and the impact of the absolute number of measurements used for the flow assessment.

Around 50 models were trained with different randomly selected experimental measurements for each undersampling percentage of measurements used. Figure 4.10 displays box plots showing the RMSE for various MFGP models' predictions compared to the *validation* subsets. In these box plots, the whiskers span from the lowest to the highest values, obtained by the different 50 models at each undersampling. ρ The box itself represents the first and third quartiles, with the median shown as a line. Each

Table 4.3.: RMSE sensitivity analysis for the H25 test case across 50 model runs different amount of measurements used.

Measurements	Median	Min	Max	Std
10%	2.5E-03	1.6E-03	4.2E-03	5.7E-04
20%	8.3E-04	6.6E-04	1.3E-03	1.8E-04
33%	4.9E-04	3.9E-04	6.3E-04	5.5E-05
50%	3.5E-04	2.8E-04	4.1E-04	3.0E-05

Table 4.4.: MaxAE sensitivity analysis for the H25 test case across 50 model runs with different amount of measurements used.

Measurements	Median	Min	Max	Std
10%	1.3E-02	6.2E-03	2.1E-02	4.4E-03
20%	4.7E-03	2.7E-03	7.1E-03	2.6E-03
33%	2.4E-03	1.6E-03	4.4E-03	6.6E-04
50%	1.6E-03	1.1E-03	2.1E-03	3.7E-04

box corresponds to a different percentage of measurements used, extracted from the full dataset test.

For a more comprehensive understanding of the RMSE results, Table 4.3 summarizes key statistics, including the median, minimum, maximum, and standard deviation of RMSE values. These metrics provide valuable insights into how different percentages of measurements affect the accuracy of flow predictions.

As anticipated, an increase in the number of measurements leads to a reduction in RMSE. This reduction is particularly pronounced when increasing the number of measurements beyond 10%. This is evident from the y-axis log scale in Figure 4.10, and is supported by the table, as median, maximum, and minimum RMSE values tend to decrease by approximately one order of magnitude. This behaviour is due to the convergence of the hyperparameters, as discussed earlier. Once a sufficient number of measurements is used, the MFGP model effectively learns the correct data hyperparameters, and additional data does not significantly improve the model.

When examining the variation among different model runs for each percentage of measurements, the robustness of the MFGP modelling approach to the location of measurements becomes apparent. As each model is tested multiple times with hyperparameter optimization restarted in an attempt to avoid local optima, the spread between the box plot whiskers can be primarily attributed to the influence of the randomly selected experimental measurements. This spread also tends to decrease with an

increase in measurements.

This observation is further supported when analysing the MaxAE sensitivity analysis results provided in Table 4.4 for a more comprehensive perspective. The MFGP model becoming more robust makes sense considering that with increasing data the chances of having a lack of sampling in relevant flow features such as the strong gradient wake region observed before are reduced and thus the inferred hyperparameters model all flow features.


In summary, the analysis highlights a clear trade-off between the number of measurements used and the accuracy of predictions. Lower percentages of measurements result in higher RMSE values, indicating less accurate predictions. This analysis underscores the challenge of achieving high prediction accuracy while minimizing the number of measurements, a main objective of the proposed methodology. It also sets the stage for the next section, which explores the use of SVGP on CFD RANS for a DoE framework in the proposed methodology. Indeed looking at the Table 4.4, the best models with 20% measurements (i.e with the minimum error), for example, present a performance that is comparable to an average model result with 33% measurements. Thus if a DoE framework was planned beforehand to be applied with the proposed methodology, an even bigger instrumentation and testing time reduction could be achieved.

Pressure Field Uncertainty Analysis


Before evaluating the proposed DoE, it is important to analyse the uncertainty in the flow predictions generated by the MFGP model. This analysis serves a dual purpose. First, it showcases the capability of the Bayesian modelling approach to predict mean flow and its associated uncertainty. Second, it provides insights into the areas in the flow field where predictions may be less certain.

Figure 4.11b displays the predicted uncertainty in total pressure ratio for the MFGP model, previously presented in Figure 4.9c and replicated in Figure 4.11a, with 33% of the measurements used. This predictive uncertainty map highlights regions within the flow field where the model might have reduced confidence in its predictions. Understanding these areas of uncertainty suggests where additional measurements or model refinement might be required to enhance prediction reliability.

The model's predicted uncertainty tends to increase rapidly outside of the studied domain, particularly evident around pitch close to zero. Furthermore, the model exhibits higher uncertainty in regions where no experimental measurements are available. However, the level of uncertainty varies across different flow regions. Notably, the highest uncertainty within the domain of interest is observed in the strong gradient wake



(a) Total pressure flow.



(b) Uncertainty.

Figure 4.11. Total pressure ratio uncertainty prediction for MFGP model presented in Figure 4.9c with 33% measurements used.




Figure 4.12. MFGP total temperature flow assessment at MP4 (right) with 33% experimental measurements undersampling against fully sampled experimental reference (left) with $\rho = 0$.

region, primarily around a radial span of 0.4, which corresponds to the location of the wake core. It is important to emphasize that even in this "high uncertainty" region, the predicted uncertainty is still within the bounds of the full experimental test's uncertainty referenced in red in the uncertainty colorbar. In other parts of the predicted flow, the uncertainty is comparatively lower.

Temperature Field Assessment

Following an in-depth analysis of the pressure field and the followed sensitivity analysis, it is relevant to perform a similar analysis for the temperature flow field. Demonstrating the capabilities of the MFGP modelling framework in predicting the temperature field underscores the versatility of this modelling approach across different flow measurements.

Figure 4.12 provides a visual comparison between the full experimental reference (on the left) and the MFGP model's predicted total temperature field with 33% of the original measurements used. Once again the white crosses represent the experimental measurement locations used in the MFGP model.

A detailed examination of the temperature field reconstruction in Figure 4.12 reveals a close visual match between the two flow fields. This result emphasizes that the MFGP modelling approach employed in the proposed methodology is not limited to total pressure measurements. The primary difference between the experimental reference and the prediction lies in the smoothness of the flow.

The fluid heats up at the hub due the rotating machinery, while at the casing, the flow exchanges heat with the atmosphere. These effect are inherent in an experimental environment but not explicitly modelled in the CFD RANS simulations. The MFGP model recognizes this pattern in the experimental data and essentially disregards the RANS results during hyperparameter optimization, as evidenced by a scaling factor (ρ) equal to 0.

The temperature field estimated RMSE is approximately 0.02%, with a MaxAE of 0.2%. These values are both below the estimated uncertainty of 0.6 K, which further supports the selection of the MFGP modelling approach within the proposed methodology. A sensitivity analysis, while not presented here for brevity, aligns with the analysis performed for total pressure.

Conclusions

In this section, the application of the MFGP framework, the main block of the proposed methodology, was applied to the H25 test case to evaluate its capabilities as a data fusion algorithm.

A consistent match was obtained on both the total pressure and total temperature flow fields in terms of critical flow features and their magnitude with a crucial note that a trade-off relation between the number of high-fidelity measurements used and the accuracy of mean flow prediction was observed. This highlighted that the model robustness is closely linked to the number of measurements, with high-fidelity measurements serving as the foundation for accurate flow predictions. The key insight is that a careful balance between the number and location of high-fidelity measurements is essential to achieve accurate predictions while minimizing testing resources.

In summary, the initial exploration into the MFGP framework confirms its robustness and adaptability for predicting flow fields. The trade-off between measurement quantity and prediction accuracy is evident, providing a solid foundation for the proposed methodology, which seeks to achieve precise flow assessments with minimal measurements. The MFGP approach potential to capture super-resolution details and its ability to account for uncertainty further validate its central role in our methodology for flow field assessment.

4.4.3. *A priori* Design of Experiments

The results above demonstrate the MFGP effectiveness in fusing high and low-fidelity data for accurate flow predictions. Another important aspect of the proposed methodol-

ogy is the DoE framework which must guide the selection of optimal sampling locations beforehand.

In this section, the results obtained using the proposed *a priori* DoE strategy with Stochastic Variational GP in the H25 compressor test case are presented. As outlined in section 3.4.1, SVGP serves as an *a priori* method for systematically identifying the most informative measurement locations based on an assumed available CFD simulation.

The SVGP DoE process begins by defining an appropriate batchsize, which was set to approximately 20% of the mesh size at MP4. This selection is made to obtain a behaviour similar to what was presented in Figure 3.6. It ensures that the estimated gradient represents the entire flow gradient adequately and the ELBO optimization converges.

Selecting Number of Measurements

Next, a crucial step in the context of the proposed methodology is the determination of the number of measurements required. For this purpose, an iterative procedure over the number of "inducing locations" (N_z) until the ELBO converges is performed. The ELBO is the equivalent of the log marginal likelihood, if it does not change, more data does not give more information to the SVGP model and the full dataset is summarized. Its optimization provides the optimal measurement locations based on the CFD simulation.

Figure 4.13 depicts the ELBO convergence for different numbers of "inducing locations," ranging from 95 to 420 (equivalent to 10% to around 45% of the full experimental database size). The ELBO tends to decrease as the number of "inducing locations" increases, reaching approximately -30000 with $N_z = 313$ (corresponding to 33% of the full experimental dataset size). Notably, further increasing the number of "inducing locations," as illustrated with $N_z = 420$, results in the same ELBO value, indicating that additional data does not significantly enhance the SVGP model's information. Thus, around 300 "inducing locations" are sufficient to effectively summarize the numerical CFD flow. A final SVGP model sets the amount of measurements and their optimal sampling locations beforehand in the proposed methodology.

SVGP Sensitivity Analysis

The experimental data was already fully acquired before this work. This means that it is not possible to directly evaluate the SVGP DoE approach, since the obtained optimal sampling locations result in a sparse grid, in contrast to the uniform grid used to acquire




Figure 4.13. ELBO optimization for hyperparameter and 'inducing locations' selection for the H25 test case CFD RANS simulation. A cluster of the ELBO from around $N_z = 300$ data points sets the DoE number of measurements required.

the H25 reference flow.

To simulate the proposed DoE with the measured dataset, an initial approach was to apply a nearest neighbour (NN) approach. The NN approach selects the experimental measured location closest to a proposed optimal sampling location. However, this leads to changes in the exact location of the proposed measurements and, in high-gradient regions, where the DoE approach proposes multiple locations, the NN approach maps them to the same available experimental point. This not only influences the DoE sampling locations but also reduces the number of measurements used in the MFGP flow assessment.

Figure 4.14 provides a closer look at the different grids in the wake region, with the SVGP selected locations and the NN approximate locations used, in green and red, respectively. The black dots represent the available measurements from the full dataset. The blue circles in the figure highlight the effect of changing the exact location of the proposed measurements, while the red boxes show the detrimental effect of merging measurement points.

As a consequence, to fully evaluate the performance of the DoE framework in the H25 reference test case, an interpolation of the full dataset was used to estimate the sup-




Figure 4.14. Nearest neighbour effect on the DoE optimal sampling grid. SVGP selected locations and the NN approximate location used, in green and red respectively.

posed measured values. To ensure accuracy in the interpolation and minimize potential errors, a single-fidelity GP model was trained using the entire experimental dataset, with the likelihood data uncertainty (σ_y) set to 0, thereby ensuring that the measured data remains unchanged. This model was then employed to predict the 'experimental measurement' at the proposed SVGP optimal locations and to conduct a similar sensitivity analysis to the one presented earlier. Additionally, this enables the evaluation of the methodology's error metrics in the validation subset, as the complete dataset can be used for assessment.

Figure 4.15 provides a box-plot comparison between random sampling, the *a priori* SVGP DoE proposed sampling approach with NN, and the *a priori* SVGP DoE proposed sampling approach without NN (and with interpolation), presented in green.

A clear improvement is obtained with the DoE approach, resulting in a significant reduction in RMSE. This reduction is not only reflected in the median but also in the reduced spread among different trained models.

While, for 10% of measurements used, the results obtained with the NN approach and the interpolation approach are similar, this does not apply to other measurement per-




Figure 4.15. RMSE comparison between random sampling, the *a priori* SVGP DoE proposed sampling approach with NN and without NN across 50 model runs for different amount of measurements used.

centages. For 20%, the median obtained is approximately the minimum value obtained with random sampling of the available experimental data.

This becomes even more evident for 33% measurements used, the amount suggested by the ELBO evaluation of Figure 4.13, where the proposed DoE, with interpolation) outperforms random sampling, achieving a similar result to the one obtained with 50% random sampling, unlike with NN initial approach. Indeed, for 50% random sampling or with the DoE approach, the box-plot result is approximately equal, and it equals the box-plot with 33% optimal measurements. This further supports the ELBO analysis performed above, showing that around 33% are the optimal number of measurements to assess this flow because adding more experimental measurements does not improve the MFGP model. The Figure suggests that the initial NN approach results in poor performance of the proposed methodology, but not the DoE framework.

To further support this analysis, Table 4.5 summarizes the obtained RMSE results for the H25 test case with the proposed *A priori* SVGP DoE. With optimal sampling and 20% of the original amount of measurements used, even for the worst obtained model, the maximum RMSE is below the estimated experimental uncertainty. Moreover, an increased robustness of the DoE approach against random sampling is proved by the RMSE standard deviation, which is one order of magnitude lower than the equivalent

Table 4.5.: RMSE sensitivity analysis with proposed *A priori* SVGP DoE across 50 different runs for a different amount of measurements used.

Measurements	Median	Min	Max	Std	Median w.r.t Random
10%	1.6E-03	1.2E-03	2.3E-03	2.5E-04	-36%
20%	6.8E-04	4.9E-04	8.6E-04	8.9E-05	-18%
33%	3.7E-04	2.9E-04	4.5E-04	4.8E-05	-25%
50%	3.2E-04	2.1E-04	4.4E-04	6.4E-05	-8.6%

in Table 4.3 for high undersampling percentages. The last column gives a quantitative value to the percentage decrease in the RMSE median value with respect to random sampling (Table 4.3) to the DoE approach, which goes as high as 36% for a reduction in the number of experimental measurements to 10%.

Conclusion

In summary, the SVGP DoE approach introduced in this section not only surpasses the performance of random sampling but also enables the application of the entire methodology for flow assessment and experimental campaign planning. The examination of different measurement percentages highlights the superiority of the DoE approach achieving a 70% measurement reduction with a 25% improved performance with respect to random sampling.

Assuming a CFD simulation is available for the operating point of interest, this approach is highly effective in an industrial context. It optimizes measurement budgets, reduces costs, and saves time by systematically selecting the most informative measurement locations, thereby streamlining the flow field assessment process for turbomachinery components.

This section results underscore the potential of the SVGP DoE framework as a valuable tool for achieving reliable, efficient, and cost-effective flow field assessment, with ramifications that extend beyond the H25 axial compressor reference test case to broader applications in turbomachinery and fluid dynamics research.

4.4.4. *In situ* Sampling

While the classic *a priori* DoE approach proposed above is an excellent choice for industrial settings where cost efficiency and careful pre-planning of experimental campaigns are paramount, the demands of academic research often extend beyond mere cost con-

siderations. In an academic research environment, different objectives might come into play. Researchers may want to focus on achieving specific uncertainty targets or probing the flow field with a higher level of detail to gain deeper insights. This is where an adaptive sampling approach can be employed in the proposed hybrid measurement technique.

The adaptive sampling approach, illustrated in Figure 3.8, commences with an initial set of strategically placed experimental measurements, achieved through an optimal Latin Hypercube Sampling (LHS) strategy. This iterative process proceeds until a predetermined stopping criterion is met. At each iteration, potential measurement locations are assessed using the Maximum Expected Predicted Error (MEPE) acquisition function. The location with the highest Expected Predicted Error (EPE) value is then selected as the next measurement point. The process is reiterated, with each new measurement added to the dataset and a new Gaussian Process (GP) model fitted. The flow field is assessed via the Multi-Fidelity Gaussian Process (MFGP) once the stopping criterion is satisfied.

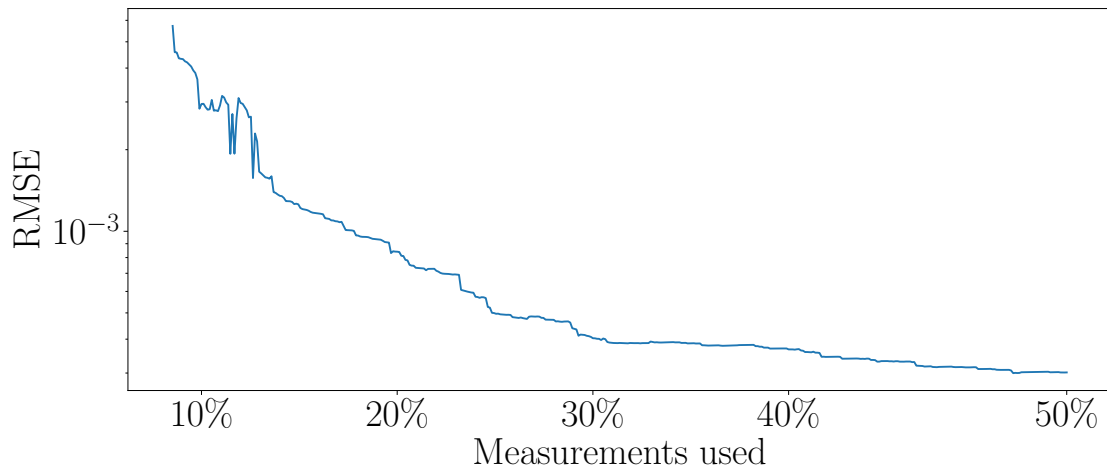
MEPE Acquisition Function Behaviour

The MEPE acquisition function is a critical component of the adaptive sampling approach. It balances the trade-off between the estimated local error and predicted uncertainty (exploitation and exploration), with the trade-off being weighted by the α factor, which is adaptively adjusted based on the previous iteration local error estimation.


To assess the performance of the adaptive sampling approach and understand how the number of measurements affect the results, a simulated run is conducted. It begins with an initial dataset size of $N_{y_{m0}} = 40$ measurements. These initial locations (\mathbf{x}_{m0}) are determined using the optimal LHS algorithm proposed by JIN et al., 2005.

The iterative process continues with the selection of measurement locations. A fixed number of iterations is defined to correspond to approximately 50% of the full dataset, this can be interpreted as the stopping criteria in this initial analysis. As measurements are added, the RMSE between the GP model predictions and the *validation* dataset, the evolution of the GP model maximum predicted uncertainty, and the α parameter are observed in Figure 4.16.


Several important observations are made. The RMSE continuously decreases as more measurements are acquired, and the maximum uncertainty in the GP model also decreases. This reduction is noticeable until around 30% of the measurements, after which it flattens. This suggests convergence in the sense that additional measurements lead to small, incremental reductions in model uncertainty.



(a) RMSE



(b) Maximum Uncertainty



(c) Weight criteria α .

Figure 4.16. Proposed adaptive sampling performance evaluation and error/uncertainty weight criteria for pressure ratio flow field assessment with increasing number of measurements.

The α parameter behaviour provides insight into different phases of the adaptive sampling approach. Initially, it tends to values close to 1. This indicates that the adaptive sampling process recognizes the need to reduce high RMSE values and prioritize exploitation. After around 15% of measurements, α tends towards values close to zero, emphasizing exploration of high uncertainty areas. At this stage, the RMSE improvement flattens, but the maximum uncertainty continues to decrease, except in specific iterations, where the adaptive sampling identifies high local error regions that lead to a direct step-like decrease in the RMSE.

A constant alternation between exploitation and exploration is seen when around 30% of the original dataset is used. However, this has little effect on the RMSE and maximum uncertainty, which are already low and comparable to the estimated experimental uncertainty.

These findings demonstrate the adaptive sampling approach ability to systematically select measurements to improve RMSE and reduce uncertainty. The adaptability of the α parameter plays a key role in efficiently guiding this process by striking a balance between exploitation and exploration.

In parallel to this analysis of the adaptive sampling approach, Figure 4.17 provides a visual comparison of this methodology at various stages of measurement acquisition, corresponding to 5%, 10%, 20%, and 33% of the measurements used for assessing the total pressure field of the H25 reference test case.

In each subfigure, the following visualizations are presented for the assessment of the total pressure field. On the top left, the reference experimental flow field with white crosses representing the experimental data points that have been "acquired". On the top right, the GP model's mean flow prediction. On the bottom left, the absolute difference between the experimental reference and the predicted pressure ratio flow and on the bottom right, the contour plot displaying the predicted GP model flow uncertainty. The RMSE, from Figure 4.16a is added to the caption.

In the initial phase (5% of measurements used), as seen in Figure 4.17a, the GP model's mean flow prediction is poor and lacks coherence in identifying flow structures, except for two high-pressure regions. At this stage, the absolute difference between the experimental reference and predicted flow (pressure ratio) is substantial, indicating a high error. This is supported by the predicted uncertainty, which is high throughout the domain, with a particular focus on the domain limits.


Both of the just mentioned effects motivate the α parameter behaviour in the phase identified around 15% measurements, with an iteration depicted in Figure 4.17b. Looking at the experimental reference and paying attention to the experimental sampling

locations, the adaptive sampling approach prioritized an evaluation of the domain limits, both in the radial and pitch coordinates.



This leads to a correct modelling of the hub and casing end wall flows, but more importantly leads to a decrease of the predicted uncertainty at the domain limits, with the contour plot showing sparse local regions of uncertainty that, at this point, approach the experimental estimated uncertainty and do not match regions of strong gradients.

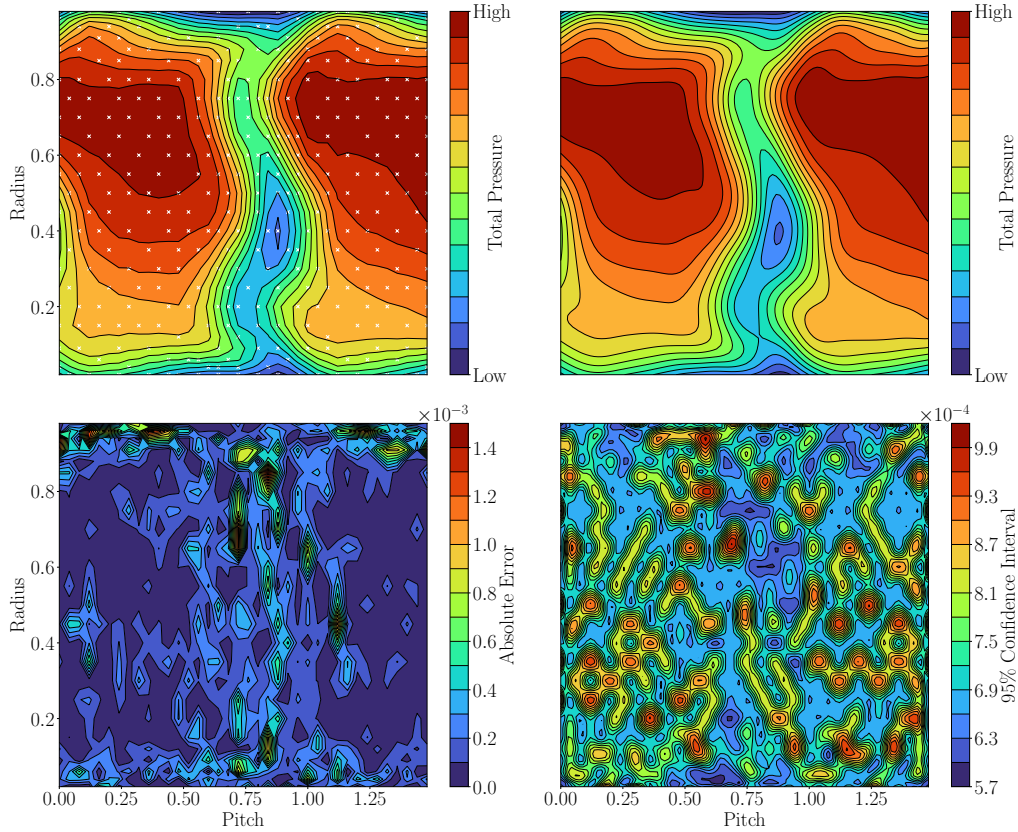
Additionally, some measurement points added during this stage are used for exploitation, targeting the wake. From 5% to 15% measurements, some measurement points are added in the gradient region of the wake that give a correct assessment of its size. However, the true depth of the wake is not correctly detected and the two high pressure regions that are periodic in the experimental reference due to representing two identical blade passages, are still not correctly identified by the adaptive sampling approach.

At 20% measurements, iteration depicted at Figure 4.17c, The GP model mean flow prediction aligns closely with the experimental reference, capturing the core of the wake and expected flow features. Looking at the absolute error contour, sparse local error



(a) 5% measurements used. RMSE $\approx 5.8E-03$.

(b) 15% measurements used. RMSE $\approx 1.2\text{E-}03$.(c) 20% measurements used. RMSE $\approx 8.8\text{E-}04$.



(d) 33% measurements used. RMSE $\approx 3.9\text{E-}04$.

Figure 4.17. Pressure flow field assessment with the proposed adaptive sampling DoE for 5%, 10%, 20% and 33% measurements used. On the top left, the reference experimental flow and the measurement samples used are shown. On the top right, the GP model mean flow prediction is presented. The bottom left shows the absolute difference between the experimental reference and the predicted pressure ratio flow. The bottom right contour displays the predicted GP model flow uncertainty.

regions are still present, however, this is due to the grid of the GP model prediction that smoothens the gradients in regions where the experimental sampling discretization cannot match the gradient size. This phase is marked by an adaptive sampling approach mainly driven by exploration. While the highest predicted uncertainty points are still chosen for measurement, the GP model’s mean prediction is highly accurate, closely matching the experimental reference.

Finally, and for the sake of completeness, Figure 4.17d shows the iteration for 33% measurements used. The GP model mean flow prediction is consistent with the previous phase, showing no significant improvement. Both the absolute error and predicted un-

certainty continue to decrease, but the flow prediction remains the same. Moreover, the uncertainty is on the order of magnitude of 10^{-4} , thus below the estimated propagated experimental uncertainty of $1.33\text{E-}03$.

The presented results confirm the adaptability of the adaptive sampling approach, where the methodology is particularly effective in the early stages, driving down error and uncertainty and, subsequently, focusing on fine-tuning the model in regions with high uncertainty and error. Once the uncertainty is close to the experimental uncertainty level, additional measurements serve to refine local uncertainty without significantly impacting the mean flow prediction.

Adaptive Sampling Sensitivity Analysis

Having explored the adaptive sampling approach performance in assessing the H25 pressure field across iterations, and understanding its potential in academic research environments, it is equally important to assess how it compares to random sampling and to the proposed *a priori* SVGP DoE approach studied above.

Figure 4.18 offers a clear comparison between random sampling, the *a priori* SVGP DoE, and adaptive sampling approaches in terms of RMSE. The adaptive sampling approach consistently outperforms random sampling. Notably, for 10% and 20% of measurements used, the median RMSE tends to surpass that of the SVGP DoE approach.

To provide a quantitative perspective on the improvements, Table 4.6 compiles the RMSE data, emphasizing the percentage improvements of the adaptive sampling approach compared to random sampling and the SVGP DoE. Notably, with 20% of measurements used in the adaptive sampling framework, the RMSE matches the one obtained with random sampling employing 33% of measurements. This result underscores the effectiveness of the adaptive sampling approach, which achieves a 31% improvement over random sampling but still lags behind the SVGP DoE with a -16% difference. This result supports the flow analysis above, where with 20% measurements all flow features are well defined and inferred. Moreover, the obtained maximum error is below the estimated uncertainty.

The robustness of the adaptive sampling approach is assessed by comparing the box plots spread of Figure 4.18 and the RMSE standard deviation column of Table 4.6 with those of the random and SVGP DoE approaches, Tables 4.3 and 4.5, respectively. The results show that for all levels of measurements used, the adaptive sampling approach exhibits superior robustness compared to random sampling. For 10% and 20% measurements used, the spread of RMSE is comparable to that of the SVGP DoE. For 33%




Figure 4.18. RMSE comparison between random, *a priori* SVGP DoE and adaptive sampling approaches across 50 model runs for different amount of measurements used.

Table 4.6.: RMSE sensitivity analysis with proposed *A priori* SVGP DoE across 50 different runs for a different amount of measurements used.

Measurements	Median	Min	Max	Std	Median w.r.t Random	w.r.t SVGP
10%	1.4E-03	1.2E-03	1.7E-03	1.7E-04	-44%	-12.5%
20%	5.7E-04	4.7E-04	6.8E-04	5.5E-05	-31.3%	-16%
33%	3.7E-04	3.3E-04	4.1E-04	2.0E-05	-25%	0%
50%	2.9E-04	2.7E-04	3.2E-04	1.3E-05	-17.1%	-9.4%

and 50%, while the median RMSE remains similar, the adaptive approach demonstrates greater robustness. This robustness arises from the reduced stochastic behaviour of the adaptive approach, where model variability primarily results from the initial set of measurements and the number of optimization restarts, resulting in consistent acquisition grids, due to the MEPE acquisition function, as measurements increase, which is not the case for the previous approaches, where different models correspond to different grids.

Despite the improved performance of the adaptive sampling for flow assessment with reduced number of measurements, an important remark is that the reduced number of measurements in this case does not directly match the reduction in testing time

when compared with the other approaches. While for the *a priori* SVGP based DoE approach, the measurement grid is planned before hand, and a reduction in number of measurements directly translates to an equivalent reduction in testing time, for the adaptive sampling approach, a GP model must be trained before a new measurement is sampled. While this model training and evaluation is performed in seconds until around 300-400 measurements points, if the number of measurements points necessary is higher, the GP model training can reach the order of minutes, recall that the GP model computational complexity is N^3 , and this additional time must be taken into account when selecting which approach of the proposed methodology to apply for the test case being studied.

The results emphasize the effectiveness of the adaptive sampling approach for flow field assessment, particularly with a reduced number of measurements. It offers improvements over random sampling and performs favourably compared to the SVGP DoE, due to the additional advantage of better robustness.

Initial Dataset Influence

Having showcased the performance of the adaptive sampling with increasing set of measurements in assessing the pressure field, one question that still imposes is the influence of the initial dataset size on the performance of the approach. Understanding the influence of the initial dataset size is an important step of assessing the performance of the adaptive sampling approach. This analysis explores how the composition of the initial dataset affects the adaptive sampling process. As discussed earlier, the initial dataset plays a role in the robustness of the approach and may introduce variability due to the randomness inherent in its construction, particularly when using an optimal LHS algorithm.

To analyse the impact of the initial dataset on adaptive sampling, three different types of initial datasets are considered: an optimal LHS algorithm, the SVGP approach proposed in this work, and rake measurements. These datasets are kept consistent with 40 initial measurements, allowing for a fair and informative comparison. The optimal LHS algorithm creates a uniformly distributed grid that minimizes the maximum-minimum distance between points. The SVGP approach relies on the prior information from RANS simulations, creating a sparse initial grid. The rake measurements, available for the test case and non-random, provide a uniform grid in both directions, with 5 rakes, each with 5 measurements, distributed azimuthally in the annulus to cover one pitch and then replicated for the one pitch and a half studied in this work.

Figure 4.19 presents box-plots comparing the different initial datasets performance




Figure 4.19. RMSE comparison between different initial datasets for the proposed adaptive sampling across 50 model runs for different amount of measurements used.

within the adaptive sampling approach for various percentages of measurements used.

The results indicate that, regardless of the number of measurements used, the choice of the initial dataset has little influence on the final RMSE. In all cases, the RMSE trends are remarkably similar. This consistency highlights the robustness of the adaptive sampling approach, which converges to stable results as measurement locations are iteratively selected, even when the initial datasets exhibit different characteristics.

One noteworthy result is that using fixed rake measurements as the initial dataset introduces higher robustness into the methodology, even at 20% of measurements used. The spread of the RMSE box-plots for the LHS and SVGP approaches appears to be more a consequence of changing the initial measurement grid than a result of the selected MEPE acquisition function performance in selecting optimal measurements locations.

The findings suggest that, when implementing adaptive sampling, the choice of the initial dataset does not significantly impact the final assessment's quality. This knowledge provides flexibility in experimental setups, allowing researchers to choose the most convenient and available initial dataset for their specific case. For instance, when prior data, such as rake measurements, are accessible, they can be used to expedite the adaptive data acquisition process. Overall, this analysis adds to the understanding of how

adaptive sampling behaves under different initial conditions and reinforces the method robustness in practice.

Stopping Criteria

Having established the potential of the adaptive sampling approach in guiding the *in situ* acquisition process, a stopping criteria, that determines when to conclude an experimental test campaign in practice, is introduced. The stopping criteria proposed for an experimental campaign must take into account various factors, including the specific objectives of the experiment, available resources, and the desired level of confidence in the results.

In situations where resources are limited, such as a fixed number of measurements or a restricted experimental testing time, the stopping criteria is straightforward - halt measurements when the limit is reached. However, in such case, the *a priori* SVGP approach, presented in the previous section is recommended and thus defining relevant stopping criteria is indispensable. It ensures that the experiment continues until both a predefined uncertainty threshold is met, and the GP model's flow assessment accurately represents the flow.

The criteria, as described in Equation 4.1, is a comprehensive and robust approach for deciding when to conclude an experimental test campaign using adaptive sampling. It combines the goals of precise flow measurement and the attainment of specific uncertainty targets, guaranteeing that the experiment continues until the desired level of accuracy is reached.

$$\sigma_{GP}(\mathbf{x}) \leq \sigma_y(\mathbf{x}) \cap \max(\text{LOOCV}(\mathbf{x})) \leq 2.58\sigma_y(\mathbf{x}) \quad (4.1)$$

The criteria consists of two parts. The first part evaluates the GP model's predicted variance (σ_{GP}^2) at each domain location (\mathbf{x}). It dictates that the GP model's predicted variance at all domain locations should be lower than the estimated experimental uncertainty (σ_y). This ensures that the model's predictions are sufficiently confident and do not exhibit significant uncertainties.

The second part of the criteria centers around the maximum LOOCV error, providing an approximate measure of local error. It requires that the maximum LOOCV error at each domain location should be lower than the 99% confidence interval of the experimental uncertainty. This guarantees that the model's mean flow predictions fall within a statistically significant range of the experimental data uncertainty and estimated local error.




Figure 4.20. Pressure flow field assessment with the proposed adaptive sampling DoE with the defined stopping criteria met. On the top left, the reference experimental flow and the measurement samples used are shown. On the top right, the GP model mean flow prediction is presented. The bottom left shows the absolute difference between the experimental reference and the predicted pressure ratio flow. The bottom right contour displays the predicted GP model flow uncertainty.

Figure 4.20 exemplifies the application of the stopping criteria in practice. At this point, a total of 171 measurements have been acquired, representing 18% of the original database size. The mean flow assessment accurately represents all flow structures in terms of shape and magnitude. The chosen measurement locations are critically positioned, covering domain limits and regions with strong flow gradients, as illustrated by the secondary flow structure at the casing merging with the wake region.

In essence, the stopping criteria ensures that the GP model's predictions exhibit both confidence (low variance) and accuracy (low local error) across the entire domain. When these conditions are met, it signals that the model has effectively captured the flow characteristics with the desired level of confidence and accuracy. This is in alignment with the objectives of academic research environments, which often require specific uncertainty targets and a deeper understanding of the flow field.

The application of the stopping criteria not only helps in achieving accurate flow assessments but also ensures efficient data collection by preventing unnecessary measurements when the desired level of confidence and accuracy has been reached. This result underscores the effectiveness of the stopping criteria in guiding the adaptive sampling approach to attain the desired level of accuracy and confidence in flow measurements, even with around 20% of the original data. It demonstrates how the approach strategically selects measurement locations to capture flow features and regions with the highest uncertainty, which is highly beneficial in academic research settings.

Conclusion

The adaptive sampling approach presented in this section offers a powerful solution for addressing the unique challenges of flow field assessment in academic research environments. While traditional *a priori* DoE may be preferred in industrial settings, academic research often demands a more nuanced approach. Researchers frequently seek to achieve specific uncertainty targets or probe flow fields with a higher level of detail to gain deeper insights. The adaptive sampling methodology is tailored to meet these demands.

Adaptive sampling, a key component of the proposed hybrid measurement technique, was presented as a solution to address these academic research needs. This technique involves iteratively selecting measurement locations until a predefined stopping criterion is met. The choice of the next measurement point is guided by the MEPE acquisition function, striking a balance between exploration and exploitation. A well-defined stopping criterion, which considers both predicted variance and local error, ensures that the experiment continues until the desired level of accuracy and confidence is achieved.

This prevents unnecessary data collection, making the approach resource-efficient.

The findings illustrate that this approach offers significant advantages. It consistently outperforms random sampling up to 44%, and for low measurement scenarios, it even surpasses the performance of the *a priori* SVGP DoE approach. Additionally, adaptive sampling demonstrates remarkable robustness, making it a reliable choice in scenarios where initial data may exhibit variability.

In summary, the adaptive sampling approach, is a valuable tool for academic research environments. It empowers researchers to achieve specific uncertainty targets and probe flow fields in great detail while ensuring efficiency in data collection. Its adaptability, robustness, and ability to systematically improve the accuracy of flow assessments make it a valuable addition to the toolkit of researchers seeking deeper insights into complex flow phenomena.

Chapter 5.

Test Case - ECL5 UHBR Fan

THE proposed hybrid measurement technique is applied to the ECL5 UHBR fan in this chapter. This final test case represents a blind application of the methodology to an experimental flow field that is unknown beforehand. It serves as the validation test case of the methodology. The experimental and numerical tools are initially presented, followed by the results obtained with both sampling approaches proposed.

Overview

5.1. Introduction	100
5.2. Experimental tools	102
5.3. Numerical tools	108
5.4. Results	108
5.4.1. Preliminaries	111
5.4.2. <i>A priori</i> Design of Experiments	112
Selecting Number of Measurements	113
Pressure Field and Uncertainty Assessment	115
Kernel Influence	121
Conclusion	125
5.4.3. <i>In situ</i> Sampling	126
Experimental test results	126
Adaptive sampling simulation	132
Conclusion	138

5.1. Introduction

Selecting test cases requires a careful balance between real-world experimental needs and the availability of data for sensitivity analysis. Ideally, chosen cases should reflect critical industrial engine components, particularly those where instrumentation limitations pose constraints on accurate flow assessment.

The H25 axial compressor test case, as detailed in the preceding chapter, provided a solid foundation. Its advantage was the availability of experimental data, providing a baseline to evaluate the methodology's performance. This case effectively functioned as a benchmark.

This chapter shifts focus to the ECL5 UHBR fan, where the methodology undergoes a validation test. Here, a blind application to the ECL5 fan test case is performed without prior experimental data. This blind test campaign evaluates the methodology performance in a real facility environment where the flow characteristics are unknown beforehand.

Table 5.1.: ECL5 stage design parameters.

Parameter	Value
Diameter	508 mm
Fan blades	16
Hub-shroud ratio	0.29
Pressure ratio	1.36
Nominal speed	11000 rpm
Mass flow	36 kg/s
Isentropic efficiency	93.4%

In addition to this reason, the ECL5 is a composite UHBR fan stage that is representative of a state of the art UHBR fan industrial architecture, due to improved propulsive efficiency and reduced noise emissions. This is achieved by increasing the engine bypass ratio and reducing the fan pressure ratio. As a consequence, the fan diameter increased and its rotational speed decreases. Compensatory measures, such as the use of lightweight composite materials, are anticipated to counterbalance the augmented engine weight.

The distinctive characteristic of UHBR fans, and ECL5, lie in their low pressure ratio and speed, making these machines more sensitive to instrumentation. This marked sensitivity is increased when compared to its high-speed counterparts, as illustrated in Figure 1.2 from the introductory chapter.




Figure 5.1. ECL5 rotor blade and fan assembly adapted from SCHNEIDER et al., 2023.

Importantly, the ECL5 represents an open research test case free from industrial restrictions. This openness extends an invitation to the broader research community, facilitating a collective exploration of its design intricacies and flow physics. In doing so, it directly contributes to advancements in UHBR engine technology.

Figure 5.1 provides a visual insight into the ECL5 rotor blade and fan assembly, adapted from SCHNEIDER et al., 2023. The accompanying Table 5.1 presents key design parameters sourced from BRANDSTETTER et al., 2019, offering a comprehensive view of the ECL5 specifications.

The significance of the ECL5 for the methodology proposed in this work becomes apparent. Its low rotational speed aligns seamlessly with the research objectives of improving experimental flow assessment methodologies tailored for next-generation compressor technologies. Moreover, the added advantage of pre-planned experimental acquisition using the proposed DoE approach further solidifies the reasoning behind using ECL5 as an validation test case.

The subsequent sections will present the experimental and numerical tools leveraged to apply the methodology to the ECL5 test case. Following this, a detailed analysis of the results will be presented, showcasing the methodology’s performance in this




Figure 5.2. ECL-B3 rig schematic.

challenging, blind application scenario.

5.2. Experimental tools

The experimental assessment of the ECL5 UHBR fan stage was performed at the ECL-B3 test rig situated at Ecole Centrale Lyon. Figure 5.2 provides a schematic overview of the facility, established in collaboration with Safran Aircraft Engines to explore fan stages at a reduced scale.

The stage, powered by an electric motor with a maximum shaft power of 3000 kW at a top speed of 16000 rpm, operates in an open-loop configuration. Intake of atmospheric air is meticulously controlled through silencers, guiding the flow into an anechoic chamber within which the test section is present. This chamber, coupled with a turbulence control screen preceding the machine inlet, ensures uniform inflow conditions and mitigates large-scale turbulence. A representative engine geometry is present at the stage flow intake. Downstream of the fan stage, a throttle valve regulates the mass flow rate, and the compressed air is expelled to the atmosphere through a lengthy circular tube section and a diffuser.

A meridional view of the ECL5 test section, shown in Figure 5.3, illustrates an




Figure 5.3. ECL5 fan stage meridional view schematic.

overview of the available instrumentation distributed across various measurement planes (MP). In this schematic, rotating components are denoted in red, while stationary components are marked in blue.

Various measurement planes host steady instrumentation designed to measure total quantities: 'in': Positioned approximately 200 mm upstream of the rotor leading edge. 'RE': In the inter-row region between rotor and stator. 'SE': Downstream of the stator. Unsteady instrumentation and optical access are strategically placed in the rotor tip region. For a detailed breakdown, refer to SCHNEIDER et al., 2023.

To evaluate the machine performance, two fixed total pressure and total temperature rakes are deployed at MP 'in,' along with two traversable rakes at the stage outlet MP 'SE'. The focus, once again, narrows down to the outlet MP 'SE' due to its critical role in stage performance and the flexibility of traversing instrumentation at this location, enabling a comprehensive flow field assessment.

For the application of the proposed hybrid methodology to the ECL5 test case, two L-shaped bent five-hole probes (5HP), depicted in Figure 5.4, are strategically positioned in the transverse system available at MP 'SE'.

The L-shape design facilitates flow measurement in proximity to the trailing edge of the stator. Both probes share the same motor control in the pitchwise direction, enabling parallel acquisition. Each probe is responsible for sampling half of the radial




Figure 5.4. L-shape bent five hole probes (5HP) deployed at MP 'SE' for flow assessment based on the proposed methodology.

span of the machine.

Additionally, the probes are bent in opposite directions, allowing for deeper penetration into secondary flow structures near the walls. One probe descends to measure deeper into the hub boundary layer, while the other ascends to measure in close proximity to the machine casing. This design ensures that, even with standard mechanical safety margins, increased end-wall proximity at the hub and casing can be achieved, as showcased in Figure 5.5, with an estimate of an end-wall proximity of 1.95 mm for both hub and casing probe, which corresponds to 1.7% and 98.3% of the span direction, respectively. The systematic uncertainty of the 5HP is around 0.00178.

To accurately benchmark the proposed hybrid methodology, reference experimental measurements are conducted to establish a baseline for evaluating the error metrics. These reference tests serve as the ground truth validation dataset against which the methodology's flow assessments are compared. Ensuring a high-fidelity assessment of the flow, 775 grid points are sampled and distributed in a matrix of 25 radial points over the span and 31 pitch-wise points over one pitch, as illustrated in Figure 5.6 with black and red dots representing casing and hub probe respectively.

A sigmoid function is employed to generate the radial grid discretization, ensuring that the reference grid is finely tuned, particularly in the boundary layer regions of both end walls. This allows the capture of intricate flow features with detail.

The use of two 5HP probes in parallel, positioned in independent blade passages, leads to a post processing complexity increase due to the need to match the two probe measured flows. It involves considering the displacement between the probes, stator pitch, and facility dimensions. This matching step ensures a correct shift of independent




Figure 5.5. End-wall proximity of the 5HP for both hub and casing with measurements possible up to 1.7% and 98.3% of the span, respectively.




Figure 5.6. ECL5 experimental test grid with black and blue dots representing casing and hub probe respectively.




Figure 5.7. Reference test pressure field contour obtained with the two L-shaped 5HP in parallel with an overlap region.

probe measurements to a normalized pitch, enabling the application of the proposed methodology.

Despite these efforts, there are experimental limitations. The periodicity assumption is evaluated in a sample region around the mean radial line, where both probes overlap. Figure 5.7 illustrates the reference test pressure field contour obtained with the two L-shaped 5HP probes in parallel, showing an overlap region of 5 radial positions for the DP point at 100NM. The slight difference in wake orientation highlights a limitation in fully respecting blade periodicity, introducing a challenge in the experimental campaign.

The wake periodicity mismatch also introduces a limitation in choosing at which radial position the casing/hub probe experimental tests are merged. The non-homogeneous and imperfectly periodic nature of real flow inside an experimental rig is acknowledged, although this work does not aim to assess the validity of this assumption. In this work, an average best match radial position was selected to set the 'ground truth' reference experimental test.

In the investigation of the application of the proposed methodology to the the ECL5 UHBR fan test case, a comprehensive exploration of various operating points (OP) is conducted. The objective is to present the versatility and efficacy of the proposed hybrid

Table 5.2.: ECL5 experimental operating points (OP).

Rotational speed	Pressure Ratio (Π)
100%	Design Point (DP)
100%	Maximum Pressure (PP)
100%	Near Choke (NC)
80 %	Design Point (DP)
55%	Design Point (DP)

methodology across a spectrum of flow conditions characterized by different pressure ratios and rotational speeds.

The reason behind this multiple point evaluation is to demonstrate the adaptability of the methodology to diverse operational scenarios. By assessing the ECL5 fan under different mass flow rates and rotational speeds, the methodology ability to adapt to various flows and generate independent optimal experimental grids for each flow condition is evaluated. This approach not only adds depth to the assessment but also substantiates the robustness of the proposed methodology under varying operational regimes.

The chosen OP test matrix is shown in Table 5.2. Each point in this matrix represents a distinct combination of pressure ratios (Π) (and mass flow) and rotational speed. Three different pressure ratio levels are studied for the nominal rotational speed: the numerical design point (DP), the point of maximum total pressure ratio (PP), and the near-choke pressure ratio (NC). For the reduced rotational speeds, 55% and 80%, only the DP is evaluated.

For each OP above, a dedicated reference experimental test is carried out. The experimental cost associated with acquiring the reference data, characterized by the aforementioned grid, is substantial. It demands a significant operating time for the machine, reflecting the resource-intensive nature of conducting this type of experimental measurements. Each OP reference test takes approximately 3 hours to acquire, considering that the two probes are being used in parallel and already reducing the time by a factor of 2.

In the subsequent sections, the outcomes obtained through the proposed hybrid methodology will be systematically compared against these reference tests. This comparison will demonstrate the methodology's capability to provide accurate flow field measurements with a significantly reduced experimental effort in terms of grid points and acquisition time.

5.3. Numerical tools

The next component for implementing the proposed methodology in this work involves the integration of numerical tools. The CFD simulations of the ECL5 utilized in this study are conducted at the Laboratoire de Mécanique des Fluides et d'Acoustique (LMFA) and are not developed as part of this work.

The primary objectives of the numerical simulations are twofold: first, to capitalize on the finely discretized flow information they provide in conjunction with the sparsely sampled experimental data, as outlined in the hybrid methodology, and second, to generate the necessary experimental grids based on the proposed SVGP DoE approach for each OP under investigation.

The numerical domain inlet is positioned 20 rotor chords upstream of the fan test section, while the outlet is located 7 stator chords downstream of the stator. Similar to the H25 axial compressor, assuming periodicity, it consists solely of one rotor fan and one stator passage. A mixing plane approach is employed to establish a connection between the rotating rotor block and the non-rotating stator block, ensuring a conservative coupling of pitch wise quantities.

For mesh generation, Autogrid5 is utilized to create a structured mesh comprising approximately 4.4×10^6 points. An O-H mesh topology is applied around the blades. Notably, all non-dimensional wall distance (y^+) values are maintained below 1 at the design point.

elsA, developed by ONERA CAMBIER et al., 2013, serves as the solver for the fully turbulent compressible steady Reynolds-Averaged Navier-Stokes (RANS) equations. The $k-\omega$ Kok model is selected as the turbulence closure model.

At the numerical domain's inlet, total pressure, total enthalpy, and flow angles are imposed. At the numerical domain outlet, a static pressure condition through radial equilibrium is applied.

5.4. Results

In this preliminary phase of the results section, a detailed comparison is conducted between the pressure fields derived from reference experimental tests and those generated by CFD simulations. Once more it is emphasized that, in the subsequent sections, the reference tests are utilized for computing error metrics and are not employed in generating experimental grids. This distinction is pivotal as the proposed methodology undergoes a blind test campaign, ensuring its applicability without relying on prior

experimental data.

The comparison between the experimental reference and CFD pressure fields provides a foundational understanding of the relation between the two datasets that are then employed in the MFGP data fusion framework of the proposed methodology. Assessing pressure distributions across various operating points elucidates the consistency and potential deviations between experimental measurements and numerical predictions.

Figure 5.8 showcases the total pressure flow field comparison between the high-fidelity reference experimental dataset (left) and the low-fidelity CFD RANS numerical flow (right). The evaluation spans the nominal rotational speed for PP (5.8a), DP (5.8b) and NC (5.8c) pressure ratios. Each subfigure is accompanied by a distinct colourbar to ensure precise comparison, and the absolute values appended to the colourbar must be considered for OPs flow comparison.


The comparison reveals a generally good agreement between the two datasets, with consistent secondary flow structures across all operating points. Notably, the stage appears to be hub-critical, with hub corner separation contributing to the loss region of the flow. However, differences emerge in the core location of the corner separation, with the experimental flow exhibiting a higher span location compared to the CFD.

The blade wake is clearly visible for all assessed OPs, with a distinctive gradient signature around 60% of the pitch. While the CFD accurately predicts the wake width, its associated gradient is weaker, as the CFD tends to predict a higher total pressure in the wake, and thus under predict the losses associated with it, due to a higher mixing. As a consequence, the RANS simulations tend to over predict the pressure ratio.


An unphysical lower-pressure region at the casing for a pitch around 20% is observed in the experimental flow but not in CFD. This discrepancy can be traced to an upstream casing instrumentation induced blockage that propagates downstream of the stator. A mask is applied to this region in subsequent sections when the experimental reference test is used to evaluate methodology performance.

A similar analysis can be performed for the 55% and 80% part rotational speed OP, with the flow comparison not shown here for conciseness. As for the 100% rotational speed DP, the experimental hub corner separation core is located on a higher span location compared to the CFD. At part speed, the pressure ratio achieved, and the losses, are lower and thus the CFD predicted mixing matches the experiments leading to a higher match in the wake region than at DP.


One difference that can be observed between the flows across all OP, is the higher loading region shape in the radial direction close to the wake signature. In the CFD, it presents a smooth round transition towards both casing and hub, while in the experi-



(a) 100% rotational speed PP.



(b) 100% rotational speed DP.



(c) 100% rotational speed NC.

Figure 5.8. ECL5 Total Pressure map comparison between experimental measurements (left) and CFD simulation (right) for the OP at nominal rotational speed.

ments the transition is abrupt and not round shaped.

Despite the highlighted differences, the overall consistency in different flow fields is considerable with an average pressure ratio difference at DP of around 1%. This accuracy in matching CFD and experimental results enhances the capabilities of the multi-fidelity framework, enabling comprehensive assessments and information gain in subsequent stages of the methodology.

Analysing the evolution of flow features across different operating points provides valuable insights. For example, at PP, hub corner separation occupies around 25% of the flow area, reaching 40% of the span. With an increase in mass flow from PP to DP, the corner separation size decreases, and its influence on losses diminishes. At NC, this flow structure is smaller, but its core exhibits extremely high losses with a absolute total pressure value lower than the atmospheric inlet.

The wake shape and size in the experimental test remains consistent for all OP, being characterized with a strong gradient that separates the highly loaded flow regions. The main secondary flow structure at the casing that propagates from the tip vortex generated at the rotor shows a higher influence on the flow at NC, occupying around 20% of the span in the blade wake region and with losses comparable to those of the hub corner separation.

The choice of testing multiple OPs, each characterized by distinct flow structure size and magnitude, aims at evaluating the robustness and adaptivity of the proposed hybrid methodology, since from a DoE perspective, the methodology adjusts the experimental grid for each specific flow condition. The proposed methodology, employing the SVGP DoE approach, customizes experimental grids to suit the unique demands of each OP. This ensures the capture of diverse flow features under different operational scenarios with optimal measurements and reduced experimental testing time.

This section sets the stage for the blind test campaign, assessing the proposed methodology's intrinsic adaptability and robustness across diverse operational scenarios.

5.4.1. Preliminaries

In the subsequent sections of this chapter, the proposed methodology is applied to assess the pressure field flow in the ECL5 UHBR fan. Two distinct fidelity data sources: high-discretized CFD RANS low-fidelity data ($\mathcal{D}_{low} = (\mathbf{x}_{low}, \mathbf{y}_{low})$) and under-sampled high-fidelity experimental data ($\mathcal{D}_{high} = (\mathbf{x}_{high}, \mathbf{y}_{high})$), where $\mathcal{D}_{low} \gg \mathcal{D}_{high}$.

The outlet flow assessment involves the input location, $\mathbf{x} = [\mathbf{r}, \boldsymbol{\theta}]$, which is mapped to the total pressure at the measurement plane 'SE', $\mathbf{y} = \mathbf{P}_T$. The input domain \mathbf{x}

is normalized between 0 and 1, and the total pressure is normalized with respect to a reference atmospheric value.

In contrast to the H25 case discussed in the previous chapter, the blind test campaign nature of this application allows the proposed optimal measurements to be fully integrated into the *training* subset. This subset is utilized in the optimization process, maximizing the marginal log-likelihood of Equation 2.14. The complete experimental reference test serves as the *validation* subset, acting as the ground truth for the methodology's predicted mean flow comparison.

To effectively apply the MFGP framework, it is imperative to ensure a match between the input domains of the different fidelities. This guarantees that the scaling parameter ρ optimized in Equation 2.14 recognizes the similarity between data sources, enhancing the high-fidelity experimental data flow assessment and enabling the direct application of the proposed DoE.

However, achieving domain matching is challenging for the ECL5 UHBR fan due to inherent differences between experimental measurements and numerical simulations. The numerical simulation domain simulates only one blade passage, located at the centre of the domain. In contrast, the experimental probes are positioned in different blade passages and at different pitch percentages and require processing of their own, as mentioned earlier.

Another limitation arises from a possible incorrect axial location estimation of the probes inside the machine due to their bending and rig installation. The axial location at which the CFD simulation is assessed may not precisely match, possibly hindering the methodology when evaluating the number of measurements required to assess the flow and their optimal locations, especially in the strong flow mixing region after the stator.

In the final experimental tests of the blind campaign, the casing probe acquisition failed, and only the hub measurements could be retrieved, adding another layer of complexity to the experimental limitations.

5.4.2. *A priori* Design of Experiments

This section illustrates the application of the proposed methodology, specifically focusing on the *A priori* Design of Experiments (DoE) approach using SVGP with the CFD simulation. The objective is to strategically determine optimal measurement locations that are then acquired experimentally to achieve a precise flow assessment within the MFGP data fusion framework.

Selecting Number of Measurements

In contrast to the fixed and evenly distributed probes in the reference experimental tests, the SVGP approach offers flexibility in both the number and location of measurements. The initial step involves evaluating the required number of measurements for the ECL5 fan test case. This iterative procedure involves varying the number of "inducing locations", N_z until the ELBO converges. The ELBO serves as the log marginal likelihood equivalent, and its optimization determines the optimal measurement locations based on the CFD simulation.

Figure 5.9 displays the ELBO convergence for different numbers of "inducing locations", ranging from 100 to 600 measurement points. The ELBO tends to decrease with an increasing number of "inducing locations," reaching approximately -50000 with $N_z = 400$. Further increases, as seen with $N_z = 500$ and $N_z = 600$, result in a similar ELBO value, indicating that additional data does not significantly enhance the information provided to the SVGP model. Hence, 400 "inducing locations" suffice to effectively summarize the numerical CFD flow data. Increasing the number of "inducing locations" further eventually leads to an increase in the numerical instability of the algorithm that can only be overcome by increasing the inherent noise associated with the data.

The higher convergence 'noise' in the ELBO, compared to the H25 test case (Figure 4.13), is attributed to the ECL5 flow complexity and the selected batch size. The stochastic batch size gradient approximation introduces higher variability due to stronger gradients and increased flow complexity but still enables convergence. Notably, exceptions in the plot, such as at $N_z = 200$ around iteration 500 and $N_z = 500$ at iteration 800, are indicative of poor gradient estimations, resembling outliers in the stochastic optimization process.

An initial error in this analysis, conducted on a CFD simulation at the wrong axial plane location, emphasized the importance of an accurate CFD-experimental match. This error resulted in an underestimation of required "inducing locations," highlighting that accurate flow representation and input domain matching is detrimental even before generating an experimental grid. This CFD simulation was located downstream of the experimental acquisition plane and as the flow mixes out and gradients weaken, the flow loses complexity, reducing the estimated need for experimental measurements.

A different grid generation is conducted for each OP under consideration. The purpose is to obtain an DoE optimal distribution of measurement points, considering the specific flow characteristics associated with different OPs, described earlier in this section. The ELBO iterative procedure is applied to all OPs, and the results are summa-




Figure 5.9. ELBO optimization for hyperparameter and 'inducing locations' selection for the ECL5 test case CFD RANS simulation for DP at 100% rotational speed. A cluster of the ELBO from around $N_z = 400$ data points sets the DoE number of measurements required.

ized in Table 5.3.

The analysis of Table 5.3 reveals that the different OPs of interest necessitate a different number of optimal measurements, with $N_z = 350$ for the PP point at nominal rotational speed and the DP at both part speeds and $N_z = 450$ for the NC point at nominal speed. This variation is clear evidence that different CFD flow fields are evaluated differently with the SVGP DoE approach, underscoring the methodology adaptability.

The interpretation of the different numbers of measurements aligns with the flow field

Table 5.3.: Number of optimal 'inducing locations' N_z for all OP tested in the ECL5 fan test case.

RPM	Π	N_z
100 %	DP	400
100 %	PP	350
100 %	NC	450
80 %	DP	350
55 %	DP	350

comparison analysis and complexity. The higher number of measurements for the NC OP at nominal speed is believed to be related to the need for more measurements to capture the strong casing secondary flow, more dominant in the NC OP than in the other flow conditions.

Contrarily, the PP point at nominal rotational speed requires fewer measurements than the DP. This might seem counter-intuitive from a turbomachinery perspective, where moving towards stall increases flow complexity. However, from a machine learning perspective, more data is required to correctly evaluate the DP separation, with higher gradients and localized phenomena. In the PP flow, the separation size has increased, and there is no need for extra data in its core since it is large, and no strong gradient is present.

This evaluation of optimal measurement locations across different OPs supports the robustness of the proposed methodology. The ability to adapt the number of measurements to specific flow conditions enhances its versatility and applicability across a spectrum of operational scenarios. The different requirements for measurements underscore the usage of SVGP as a valid a priori DoE approach based on the unique demands of each OP.


Pressure Field and Uncertainty Assessment

Following the generation of grids for each OP and the experimental data acquisition, a MFGP model is trained for each OP to evaluate the full flow field in a finely discretized grid, informed by the combination of the high-fidelity measurements acquired at an optimal grid obtained from the SVGP DoE and the low-fidelity CFD data.


Figure 5.10a illustrates the mean inferred total pressure contours for the DP at design rotational speed, on the right, compared against the experimental reference test, on the left. Moreover, an integral aspect of the MFGP model is its capacity to provide uncertainty estimates for the inferred flow field. Figure 5.10b illustrates the uncertainty distribution accompanying the inferred total pressure contours.

The MFGP model, trained with a combination of high-fidelity measurements and low-fidelity CFD data, exhibits a flow pattern similar to the experimental reference, effectively capturing the hub corner separation. The core, not sampled in the experimental reference but present in the CFD, is successfully exploited by the MFGP framework.

The observed mismatch between the flows is primarily in the wake region, where the MFGP smoothens the sharp wake gradient seen in experiments. This smoothing effect is attributed to both the inherent behaviour of the MFGP, which tends to smooth



(a) Mean flow assessment (right) against fully experimental reference (left).



(b) Uncertainty.

Figure 5.10. ECL5 mean total pressure and uncertainty contours for DP at 100NM inferred with *a priori* SVGP DoE and MFGP model. Comparison to experimental reference.

function behaviour, and the CFD simulations, which introduce higher mixing than observed in the experimental reference. Another inherited trait from the CFD is the rounded pattern "bump" along the radial direction in the wake region that is not present in the experiments. If it is assumed that this "bump" is caused by a wrong estimation of the mixing by the RANS, the MFGP modelling is acquiring a flow structure that is poorly modelled by the low-fidelity data. Another possible cause for the "bump", that facilitates its inherence from the CFD, is the smoothing of the sharp gradients by the kernel selection in the radial direction.

The influence of CFD simulations on the MFGP model is evident in the scaling factor ρ , close to 1, indicating that the MFGP model perceives the CFD as a reliable data source when assessed with the experimental measurements, since the flow structures size match and this information is learned by the independent fidelities length scales. The model interprets the differences between the CFD and experiments with the local bias $\delta(\mathbf{x})$ term in the MFGP formulation of Equation 2.21.

Figure 5.10b illustrates the uncertainty distribution accompanying the inferred total pressure contours. The model, characterizing the differences between data sources as noise and exhibiting high correlation datasets ($\rho \approx 1$), adopts a conservative stance in propagating uncertainties, as expected from Equation 2.24, the high fidelity predicted uncertainty is a sum of its inherent uncertainty and the relation with the low fidelity uncertainty.

The model interpretation of different data sources as noise directly influences the estimated model uncertainty. Despite an overall high uncertainty (with an average below 0.01), specific regions where the CFD and experimental data agree show uncertainty as low as 0. However, in the hub region, where strong boundary layer gradients differ between CFD and experiments, the uncertainty peaks at around 0.0168 Pa, a value much higher than the systematic uncertainty of 180 Pa.

Due to an overall high uncertainty it is important to identify clearly regions where the model mean flow deviates from the reference test to ensure that the flow field is regardless accurately inferred. To further evaluate the difference between the reference experimental data and the inferred flow field, and identify regions where the error is high, Figure 5.11 displays the absolute pointwise pressure difference. The error is negligible except in strong gradient regions, notably in the wake, where the maximum error is around 3600 Pa.

Despite the high overall uncertainty, the mean flow field is accurately captured. The observed discrepancy in the absolute pressure difference map highlights the importance of targeted refinement strategies, supporting the proposed adaptive sampling approach.




Figure 5.11. Absolute pressure difference between reference experimental data and the inferred flow field, highlighting local error regions.


Table 5.4.: Quantitative evaluation of the relevant estimated hyperparameters and the error metrics for the MFGP models for the different ECL5 OPs.

RPM	Π	N_z	ρ	σ_{high}^2	σ_{low}^2	RMSE	MaxAE	R^2
100%	DP	400	0.993	≈ 0	2.45E-05	0.0054	0.0345	0.995
100%	PP	350	0.991	≈ 0	1.55E-05	0.0047	0.0263	0.996
100%	NC	450	0.984	≈ 0	4.68E-05	0.0076	0.0568	0.995
80%	DP	350	0.991	≈ 0	1.13E-05	0.0052	0.0281	0.988
55%	DP	350	0.998	≈ 0	4.86E-06	0.0024	0.0138	0.990


For completeness, Figures 5.12 and 5.13, show different mean inferred total pressure and uncertainty contours for the nominal speed PP and NC OP, respectively. Each individual experimental reference is also illustrated. The flow and uncertainty analysis is analogous to the one performed for the 100NM DP flow and thus shows the ability of the proposed methodology to infer the experimental flow field and estimate their uncertainty with a sparse optimally reduced number of measurements across different flow conditions.

Table 5.4 provides a quantitative evaluation of relevant estimated hyperparameters and error metrics for the MFGP model trained for each OP. This comprehensive assessment complements the qualitative analysis, demonstrating the effectiveness of the MFGP framework in inferring the flow fields across the different OPs.

The scaling factor ρ is close to 1 for all cases, suggesting that the MFGP model interprets the CFD data as a reliable source when assessed alongside experimental measurements. This indicates a good match and similarity in flow structures between




(a) Mean flow assessment (right) against fully experimental reference (left).




(b) Uncertainty.

Figure 5.12. ECL5 mean total pressure and uncertainty contours for PP at 100NM inferred with *a priori* SVGP DoE and MFGP model. Comparison to experimental reference.



(a) Mean flow assessment (right) against fully experimental reference (left).



(b) Uncertainty.

Figure 5.13. ECL5 mean total pressure and uncertainty contours for NC at 100NM inferred with *a priori* SVGP DoE and MFGP model. Comparison to experimental reference.

the CFD and experimental data.

σ_{high}^2 and σ_{low}^2 represent the variances associated with the high-fidelity and low-fidelity data, respectively. The zero value for σ_{high}^2 means that the model forces a fit to the experimental data where it is available and this sustains the zero uncertainty locations observed in the above Figures, further sustained by the observed values of R^2 , ranging from 0.988 to 0.996, that show the model captures the variance in the experimental data well.

The high value σ_{low}^2 indicates that the model attributes any differences between the high-fidelity and low-fidelity data to noise and thus, where no data is available the uncertainty increases. This conservative approach ensures that uncertainties are appropriately captured.

The error metrics, presented in terms of pressure ratio, indicate a good overall fit of the models for the different flow conditions. The NC condition shows the highest predicted error, and the highest uncertainty, with a RMSE of around 750 Pa. This OP, for which more experimental measurements N_z were also required, appears to be more complex flow from a modelling perspective.

The MaxAE values for all models are one order of magnitude higher than the RMSE, with a range from 0.0138 (at part speed design) to 0.0568 (at NC). This suggests that these high errors are present only in point wise regions of high gradients that are poorly sampled with the experimental discretization.

In summary, the quantitative evaluation supports the qualitative assessment of the MFGP framework effectiveness in capturing flow features and providing accurate predictions across different operating points of the ECL5 UHBR fan. The adaptability of the proposed methodology to different flow conditions is clearly addressed, and the conservative approach to uncertainty propagation, coupled with an accurate mean flow inference enhances the reliability and validates the proposed methodology in a blind test case.

Kernel Influence

Since the initial description of the proposed methodology, the Matern 5/2 kernel, introduced in equation 2.13 has been employed in all GP models. This choice was motivated by the kernel documented performance in the literature and its simplicity, making it easy to implement and contributing to the overall methodological efficiency.

Following the comprehensive pressure field and uncertainty assessment detailed above, a focused study on kernel influence is conducted. The objective is to refine the assessment of the flow against the experimental reference test, especially in the wake region.

This study does not aim to undermine the positive outcomes of the previous results but rather explores if alternative, more intricate kernels can enhance flow assessment at the expense of increased complexity and computational cost. To mitigate this cost, the subsequent results are derived using a single-fidelity GP model.

Two new kernels are introduced and applied to the same experimental test for the ECL5 fan, specifically the DP at nominal rotational speed.

The first introduced kernel is the change point (CP) kernel, designed to facilitate an abrupt transition between two distinct kernels. This choice is made with the anticipation that it can effectively capture the pronounced wake gradient in the flow. Mathematically, the CP kernel is defined as:

$$CP(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{K}_1(\mathbf{x}_i, \mathbf{x}_j)(1 - \sigma(\mathbf{x}_i))(1 - \sigma(\mathbf{x}_j)) + \mathbf{K}_2(\mathbf{x}_i, \mathbf{x}_j)\sigma(\mathbf{x}_i)\sigma(\mathbf{x}_j) \quad (5.1)$$

where \mathbf{K}_1 and \mathbf{K}_2 are distinct kernels, both chosen here as independent Matern 5/2 kernels. The sigmoid function, $\sigma(\mathbf{x})$, controls the smooth transition between the two kernels, with parameters s determining the smoothness of the step and x_0 indicating the domain location at which the transition occurs:

$$\sigma(\mathbf{x}) = \frac{1}{1 + e^{-s(x-x_0)}} \quad (5.2)$$

These parameters, in addition to the new Matern 5/2 kernel are considered hyperparameters to be inferred with the marginal log-likelihood optimization. For this application, the CP kernel is selectively applied in the pitch-wise direction to model the wake gradient.

The second introduced kernel is a custom composite, crafted from insights derived from CFD simulations and the underlying flow physics. It is a combination of multiple individual kernels, each tailored to model specific flow behaviours observed in the simulation. The kernel composition includes:

Table 5.5.: Comparison of error metrics between the different kernels for the pressure field inference of the ECL5 fan at DP for 100% rpm.

Kernel	RMSE	MaxAE
Matérn 5/2	0.0054	0.0345
Change Point	0.0049	0.0297
Custom	0.0046	0.0316

$$\mathbf{K}(\mathbf{x}_i, \mathbf{x}_j) = \sigma_{f1}^2 \exp\left(-\frac{\sin(\pi|\mathbf{x}_i, \mathbf{x}_j|/P)}{2\ell_P}\right)^2 + \sigma_{f2}^2 \left(1 + \frac{(\mathbf{x}_i - \mathbf{x}_j)^T (\mathbf{x}_i, \mathbf{x}_j)}{2\ell_{RQ}^2}\right)^{-\alpha} \exp\left(-\frac{(\mathbf{x}_i - \mathbf{x}_j)^T (\mathbf{x}_i, \mathbf{x}_j)}{2\ell_{RBF}^2}\right) \quad (5.3)$$

This kernel accounts for periodic blade passage behaviour, local periodicity within the blade passage, and continuous regions of high loading in both flow directions. The rational quadratic kernel is incorporated to model smooth variations across multiple length scales, while the squared exponential kernel facilitates smooth local changes in the radial direction. The parameters σ_{f1} , ℓ_P , P , σ_{f2} , ℓ_{RQ} , α and ℓ_{RBF} are hyper-parameters to be inferred through marginal log-likelihood optimization, considerably amplifying model complexity.

Table 5.5 presents a comparative analysis of the RMSE and MaxAE error metrics for pressure field inference in the ECL5 fan at DP with 100% RPM using the different described kernels.

The CP and custom kernels demonstrate lower RMSE compared to the Matérn 5/2 Kernel. This indicates that these kernels, designed to capture specific flow characteristics, provide around 15% more accurate overall representation of the pressure field. This better representation is highlighted in the MaxAE, with a substantial decrease of around 500 Pa. This implies that these kernels perform better in capturing the modelled localized flow phenomena, especially in regions with high gradients like the wake region, where the error was higher with the Matern kernel, as depicted in Figure 5.11. This is visually supported by Figure 5.14 that compares the different kernel models inferred pressure ratios at a fixed radius of $r = 0.63$ against the experimental reference. It is clear that both the CP and the custom kernel outperform the Matern52 kernel

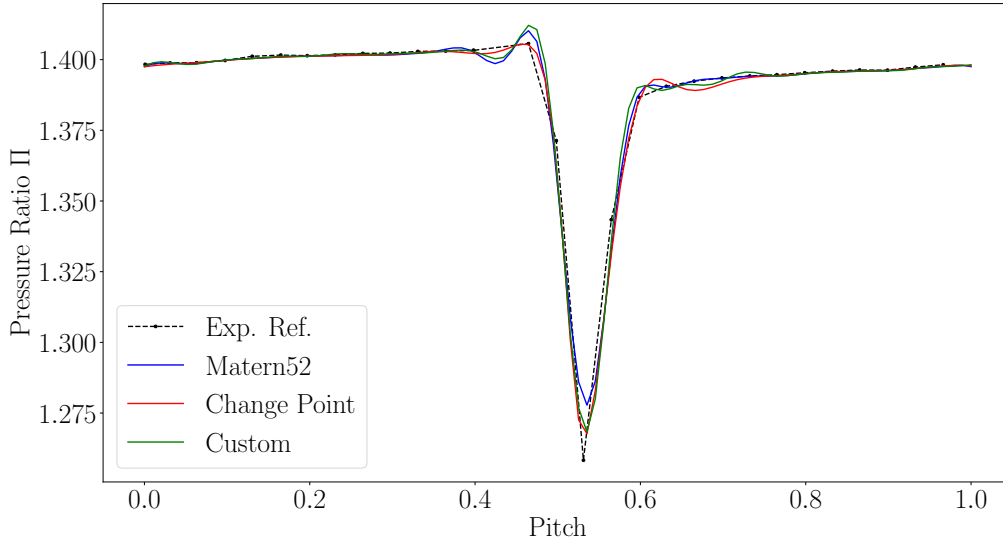


Figure 5.14. Pressure ratio pitch comparison of the different kernel at for a fixed $r = 0.63$.

in the inference of the wake core and that the CP kernel models better the gradient transition from the clean flow to the wake.

The observed improvement in accuracy with the CP and custom-built kernels introduces a trade-off, especially in terms of model complexity. Occam's Razor, a principle of problem solving, suggests that among competing hypotheses, the one with the fewest assumptions should be selected. In the context of GP modelling, this translates to preferring simpler kernels when they provide comparable performance.

The Matérn 5/2 kernel is a simpler choice, with fewer hyperparameters, and might be preferred when computational efficiency is critical. Its application aligns with Occam's Razor, favouring simplicity unless a more complex model significantly improves performance. Moreover, with the introduction of additional hyperparameters, the models become more intricate. While they demonstrate enhanced accuracy in this specific scenario, there is a risk of over fitting. The increased complexity of these kernels raises the possibility of tailoring the Bayesian prior of the model too closely to the training data, leading to reduced generalization performance on unmeasured flow locations.

Another important remark is that, unlike the Matérn 5/2 kernel, which can be applied in a more "black box" manner, the CP and custom-built kernels are proposed with expert knowledge of the GP modelling framework and a deep understanding of the physical processes governing the system. This approach aligns with the principle of incorporating physics knowledge into machine learning models, ensuring that the

additional complexity is purposeful and theoretically grounded. This approach is particularly valuable in situations where the physics of the system are well-understood and can be translated into the kernel structure.

The expert knowledge is inherited by trained kernel hyperparameters. In the CP kernel, the change between kernels occurs at $x_0 = 0.52$, corresponding to the physical wake location. For the custom-built kernel, the detection of $P \approx 1$ corresponds to the expected periodicity of the blade passage, reflecting the incorporation of domain knowledge into the model.

In summary, the CP and custom-built kernels show promise in enhancing the accuracy of pressure field inference. However, the decision to adopt these kernels should be based on a comprehensive evaluation of their benefits against the associated increase in model complexity and user required input knowledge.

Conclusion

This subsection presents a detailed analysis of the results obtained from the application of the *a priori* DoE strategy in the proposed methodology to the ECL5 UHBR fan blind test campaign.

The *A priori* DoE phase, employing the SVGP approach, demonstrates the adaptability of the proposed methodology. Different OPs require varying numbers of optimal measurements, reflecting the unique demands of each flow condition. This adaptability ensures the capture of diverse flow features under different operational scenarios with optimal measurements. Nevertheless, a reduction on the number of measurements of 50% was obtained, thus reducing a complete map assessment from 3 to around 2 hours.

The MFGP framework is then applied to assess the pressure field flow in the ECL5 UHBR fan. The mean inferred total pressure contours demonstrate the ability of the model to accurately capture flow features, including hub corner separation, despite discrepancies in the wake region, but with a pressure ratio maximum RMSE of 0.007 against the reference test, throughout all operating points.

The uncertainty assessment reveals the model conservative stance, due to a high dependence on the RANS simulation. The error metrics, confirm the effectiveness of the MFGP framework across different OPs, despite the higher uncertainty than the measurements systematic uncertainty.

A focused study on kernel influence introduces two new kernels, the CP kernel and a custom composite kernel. These kernels demonstrate enhanced accuracy in capturing localized flow phenomena, particularly in regions with high gradients like the wake. However, the decision to adopt these kernels should consider the trade-off between

increased accuracy and higher model complexity.

In conclusion, the proposed methodology exhibits promising results in its adaptability, robustness, and accuracy across different operational scenarios of the ECL5 UHBR fan. The ability to integrate prior knowledge into the GP model, as demonstrated by the choice of kernels, further enhances the reliability of the methodology.

The successful application of the proposed approach in a blind test case underscores its potential for efficient and accurate aerodynamic assessments in turbomachinery design.

5.4.3. *In situ* Sampling

In this section, the results obtained by applying the adaptive *in situ* sampling approach of the proposed methodology to the ECL5 fan test case in an experimental blind test campaign are presented. The goal is to analyse the experimental measurements' locations and the adaptability of the sampling approach to recognize flow structures and focus the sampling efforts in correctly assessing the flow while considering the predicted modelling uncertainty.

Experimental test results

Due to the dual probe parallel acquisition system, during the iterative procedure, to avoid wasting a possible measurement, both probes were used to acquire data points. Recall that the experimental probes are controlled by the same motor in the tangential direction. This means that after the MEPE acquisition function evaluates the next optimal location, both probes are translated simultaneously to that pitch location. However, each probe has an independent radial motor control. Thus, a second evaluation of the MEPE acquisition function is performed, limited to the pitch of the first optimal measurement and to the span of the non-optimal probe. For example, if the MEPE acquisition function determines that the next optimal measurement location is at a pitch of 0.8 at a radius of 0.9, close to the casing, the second evaluation limits the candidate points to the pitch of 0.8 and to the radial span that the hub probe can measure.

The analysis starts by presenting the raw data obtained from the test. Due to time limitations in the test rig, only the DP at 100% rotational speed was investigated.

Following the discussion on the initial dataset influence in the previous chapter, an initial set of $N_{\mathbf{y}_{m0}} = 40$ measurements were taken with an optimal Latin Hypercube Sampling (LHS), with each individual probe sampling half of them.




Figure 5.15. Initial set of $N_{y_{m0}} = 40$ measurements obtained with optimal LHS for the ECL5 test case.

Figure 5.15 shows the obtained grid and the measured point-wise total pressure ratio. As expected from the LHS algorithm, the measurements are uniformly distributed in the entire domain. The measured pressure ratio values range from 1.15 to 1.4, similar to what is measured in the reference test in Figure 5.8b. Additionally to this evidence, with the LHS grid, the signature of secondary flow structures such as the wake and the hub corner separation are present in the flow, and the probes accurately measure.

The adaptive sampling blind test continued until a total of 220 measurements were acquired. The obtained measurement grid and the measured pressure values are presented in Figure 5.16.

The adaptive sampling process started by measuring the domain limits on both hub and casing correctly, demonstrated by the low pressure values measured. However, around mid-span, all the casing probe measurements present the same total pressure measured. Adding to this, there is a cluster of measurements points in the middle of the pitch passage, where the wake should be measured, that measure a high and constant pressure.

This region was correctly measured with the initial set of measurements, as noted in the analysis of Figure 5.15 above, and it is believed that the observed behaviour is due to a problem that occurred during the adaptive sampling acquisition/processing of




Figure 5.16. Final $N_y = 220$ measurements obtained with the adaptive sampling blind test for the ECL5 test case.


the casing probe. It led to a continuous evaluation of a constant pressure at a constant radius, which gave contrary information to the adaptive sampling method. From the initial set of measurements, it expects a low pressure in this region, causing the heavy sampling verified in the region.

To address the paradoxical data measurements in the mid-span region with the casing probe, the GP model justifies the data by inferring that a large amount of noise is present in the measurements, eventually leading to the termination of the test.


Despite that, the hub probe measurements seem consistent with what is expected of the flow field. Moreover, looking at the adaptive sampling grid of the hub probe, there is a clear signature of the hub corner separation, with the adaptive sampling focusing measurement locations on this region.

With the goal of exploiting the correct data acquisition of the hub probe, the measurements that correspond to it are used to build a GP model and evaluate the lower radius section of the flow. Figure 5.17 shows the obtained pressure result with the GP model during the adaptive sampling acquisition for the ECL5 blind test case for the lower radius. The pressure field assessment is shown in Figure 5.17a, and the corresponding uncertainty is shown in Figure 5.17b.

The mean flow reconstruction shows a clear signature of the hub corner separation.



(a) Mean flow assessment.



(b) Uncertainty evaluation.

Figure 5.17. ECL5 lower radius pressure field and uncertainty assessed with the GP model with the adaptive sampling measurements of the hub probe.

Moreover, its size and magnitude match with the ones measured in the experimental reference test, with the different being that the GP model results in a smooth flow due to the ability to assess the flow in a finely discretized grid.

However, the flow does not have enough measurements, and the wake signature is not clearly defined since the probe only sampled the hub corner separation region, and the measurements around mid-channel high are located in the high-pressure region. Despite that, the GP model is fully aware of this situation, as the predicted uncertainty contour shows a maximum uncertainty in the mid-channel high as high as 0.05, the area where the wake should be visible.


The contrast between uncertainty levels around mid-channel and at the hub is evident and comes from two sources. One is the fact that at the hub, measurements were acquired, and two, the high-loaded blade region appears in the hub probe domain region. The model is unsure if this is correctly measured since the hub flow does not share any length scales with this high-pressure region. This inconsistency in the domain evaluated and the low number of measurements leads to an overall increase of the predicted uncertainty, with the model predicting uncertainty levels of around 0.0197 even where experimental measurements are available.

The observed behaviour, although limited to the hub probe, leads to an interesting discussion. On one hand, the adaptive sampling methodology appears to correctly reconstruct the hub corner separation, with the optimal sampling locations focusing on its strong gradient. On the other hand, there is a clear lack of information on the complete flow domain that translates to poor confidence in the model's reconstruction, as evidenced by the high predicted uncertainty, even where measurements are available.


In an attempt to overcome this, the final step of the proposed methodology is to apply the MFGP framework. This allows leveraging the full domain information with the hub probe measurements to predict the complete flow domain.

Figure 5.18 shows the obtained pressure result with the MFGP model after the adaptive sampling acquisition for the ECL5 blind test case for the lower radius. The pressure field is shown in Figure 5.18a with the experimental reference on the left and the MFGP assessment on the right, and the corresponding uncertainty in Figure 5.18b.

The pressure field shows, as expected, the correct assessment of the hub corner separation size and width. The flow in the lower radial region matches the experimental reference. The wake region, on the other hand, shows a poor match, with a smoother gradient than the experimental reference. This effect is expected since, as analysed before, the CFD simulation tends to over predict the mixing of the flow in the wake region. The CFD simulation, however, was misplacing the location of the hub corner



(a) Mean flow assessment.



(b) Uncertainty evaluation.

Figure 5.18. ECL5 pressure field and uncertainty evaluated with the MFGP model after the adaptive sampling measurements of the hub probe.

separation, but this is improved with the high-fidelity experiments in this region.

The upper radius of the flow, which should have been measured with the casing probe, is inferred by leveraging the CFD simulation. Despite no experimental measurements in this region, the flow matches against the experimental reference, even in the secondary flow structures in the casing. Again, the discrepancy against the reference test is solely in the wake region near the mid-radius where the mixing is over predicted. Additionally, the shape of the high-loaded region is rounded, as observed before in the CFD. This analysis is further supported by the MFGP hyperparameters, with the scaling factor ρ close to 1.

No high-fidelity data is available in this region, and this is shown in the uncertainty contour with an average value around 0.055. An interesting analysis is to compare the uncertainty contour obtained with the MFGP model against the one obtained with the simple GP model evaluated during the adaptive sampling process.

Despite the high predicted uncertainty values, reaching 0.005 even in measurement locations, due to the low number of measurements and the not fully measured domain, the uncertainty is lower than if the multi-fidelity framework is not applied, as shown in Figure 5.17b, where the uncertainty in measurement locations is at least 0.015.

Using the MFGP model, after the adaptive sampling approach is finished, especially in this case where the number of measurements is lower than necessary, allows for inferring regions of the flow not sampled and contributes for a proper uncertainty evaluation of these regions.


Adaptive sampling simulation

To further analyse the proposed adaptive sampling and understand the impact of missing casing measurements, a simulation similar to the one applied for the H25 axial compressor test case is performed.

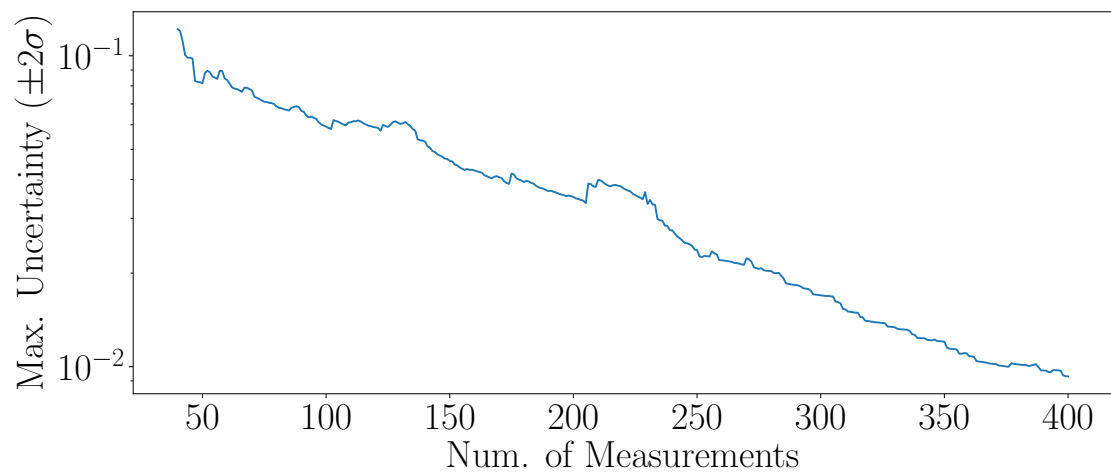
The initial set of measurements acquired in the blind test is used, and the experimental reference test is employed to generate the adaptive sampling measurements needed during the iterative procedure. The iteration process is halted with a fixed number of measurements equal to 400.

To analyse the evolution of the flow inference as measurements are added, the RMSE between the GP model predictions and the reference test, the evolution of the GP model-estimated maximum uncertainty, and the weight parameter α are illustrated in Figure 5.19.

The RMSE and the estimated maximum uncertainty decrease continuously with an increasing number of measurements.



(a) RMSE



(b) Maximum Uncertainty.


(c) Weight criteria α .

Figure 5.19. Simulated ECL5 adaptive sampling error/uncertainty and weight criteria for pressure ratio flow field assessment with increasing number of measurements.

After an initial phase with a step decrease in the metrics (around 50 measurements), the RMSE decrease follows a constant gradient until around 230 measurements. Afterward, as more measurements are added, the RMSE tends to flatten while the maximum predicted uncertainty keeps decreasing. At 400 measurements, where the iteration process was stopped, there is no conclusive convergence of the estimated maximum uncertainty but the RMSE appears to approach a limit trend.

The convergence analysis above highlights two different regions, separated by a strong reduction in both evaluated metrics. These two regions are visualized in the weight parameter α . In the initial constant decrease phase, α continuously changes between 0 and 1, balancing the trade-off between exploration and exploitation.

After the sharp RMSE decrease around 230 measurements, the behaviour of the acquisition function changes, with α mainly taking values of 0. This means that the adaptive sampling focuses mainly on exploration and decreasing the overall uncertainty in the domain. This is further supported by the flattening of the RMSE evolution described above, as more measurements are not improving the mean prediction of the flow but rather decreasing the underlying uncertainty the model estimates.

In summary, with the adaptive sampling strategy proposed, it is expected that the flow structures are correctly identified with around 230 measurements, and afterward, the model mainly focuses on sampling areas with high uncertainty.

To further support the above analysis, Figure 5.20 shows the 'simulated' pressure flow field assessment for the ECL5 fan at DP with the adaptive sampling procedure at various stages of measurement acquisition, corresponding to 120, 220, 240, and 360 measurements.

In each subfigure, the following visualizations are presented for the assessment of the total pressure field. On the top left, the reference experimental flow and the measurement locations are shown. On the top right, the GP model mean flow prediction is presented. The bottom left shows the absolute difference between the experimental reference and the predicted pressure ratio. The bottom right contour displays the propagated predicted uncertainty.

Figure 5.20a, corresponding to 120 measurements acquired, represents the initial phase of adaptive sampling, where the constant decrease of error starts. At this moment, the mean flow inference begins to recognize important flow patterns such as the low-pressure region at the hub and the higher loading region at mid-span. The recognition of these features is due to the focus of measurements in these high gradient locations that the adaptive sampling considers critical.

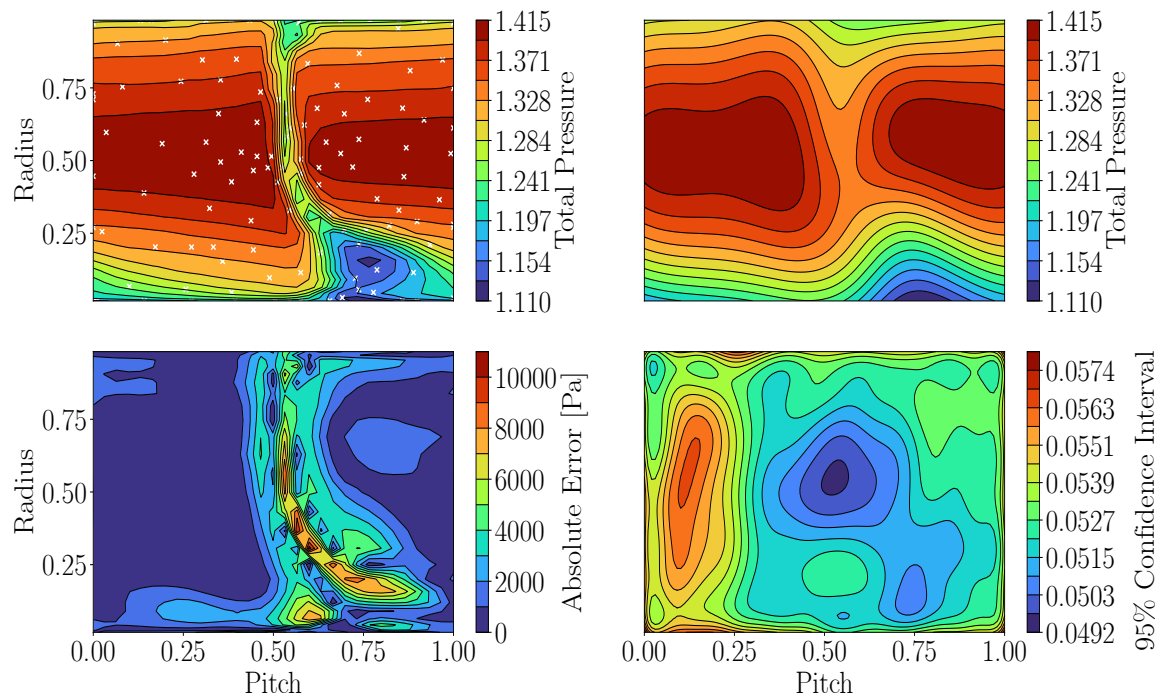
As the number of measurements increases to 220, Figure 5.20b shows that the main

characteristics of the flow are now visible. There is a separation at the hub, despite its poor prediction, and the wake signature begins to be visible, separating two higher loading regions at mid-span and connecting to the secondary flow region at the casing, despite showing an underprediction of its depth.


The experimental grid at this moment retains focus on the wake region with a collapse of measurements in the high gradient, with a consequential reduction in the error, with some iterations sampling other locations of the flow to reduce overall uncertainty, justifying the α behaviour in this region.

This trade-off between exploration and exploitation continues until around 230 measurements are acquired, and the sharp error decrease discussed above occurs.


Figure 5.20c shows the adaptive sampling inference with 240 measurements acquired, after the just mentioned sharp decrease. The mean flow reconstruction correctly identifies all flow features and matches the reference experimental test. The increased accuracy in the flow inference is attributed to the hub corner separation and the hub boundary layer being assessed as different flow structures. Moreover, the wake absolute pressure depth is evaluated correctly. This is supported by the analysis of the absolute error difference between experimental reference and inference that shows an error under 2000 Pa in all these regions. The only regions with high sparse errors are around the




(a) 120 measurements acquired. RMSE $\approx 2.2\text{E-}02$.



(b) 220 measurements acquired. RMSE $\approx 1.3\text{E-}02$.



(c) 240 measurements acquired. RMSE $\approx 8.3\text{E-}03$.



(d) 360 measurements acquired. $\text{RMSE} \approx 6.0\text{E-}03$.

Figure 5.20. Pressure flow field assessment with the proposed adaptive sampling DoE for 120, 220, 240, 360 measurements acquired. On the top left, the reference experimental flow and the measurement samples used are shown. On the top right, the GP model mean flow prediction is presented. The bottom left shows the absolute difference between the experimental reference and the predicted pressure ratio flow. The bottom right contour displays the predicted GP model flow uncertainty.

casing, with the mean flow slightly underpredicting point-wise the pressure loss and around the hub at pitch around 0.5, where the main flow interacts with the separation and boundary layer simultaneously.

From the correct inference of the mean flow, the adaptive sampling behaviour changes towards exploration, as the parameter α remains 0 for most iterations. Due to the high concentration of measurements in the wake and hub corner separation, these regions of the flow are considered the least uncertain, showing values of 0.008, which is the GP model estimated data noise uncertainty.

For completeness, Figure 5.20d shows the continuation of the adaptive sampling procedure in its fully exploration phase, with 360 measurements acquired. There is no remarkable difference in the mean flow assessment except for the improved point-wise match in the sparse error regions analysed above.

However, there is a clear decrease in the uncertainty prediction of the flow with the maximum uncertainty decreasing from 0.022 to around 0.011 and the minimum decreasing from 0.008 to 0.0012, following the decrease of the estimated data noise. Despite the similar contour plot, the high regions of uncertainty for 270 measurements become regions of low uncertainty with 360 measurements. This emphasizes again that adaptive sampling is targeting high uncertainty areas in this phase.

In summary, the simulation of the adaptive sampling strategy for the ECL5 fan test case provides valuable insights into its performance and behaviour. The results demonstrate that, with approximately 230 measurements, the proposed adaptive sampling method effectively identifies and captures key flow features, such as the hub corner separation and wake structures. Subsequent measurements beyond this point primarily contribute to reducing uncertainties rather than improving the mean flow prediction.

Conclusion

In this section, the adaptive *in situ* sampling approach as part of the proposed methodology to assess the ECL5 fan test case in an experimental blind test campaign is presented. The objective was to evaluate the performance of the sampling approach in recognizing flow structures, optimally sampling the flow in real time effectively, and thus giving a correct flow assessment.

The experimental test results revealed the complexity of the dual probe parallel acquisition system and the challenges posed by simultaneous translations. Despite encountering issues during the adaptive sampling blind test, particularly in the mid-span region with the casing probe, the methodology demonstrated success in capturing the hub corner separation with the hub probe.

To address the limitations of incomplete flow domain information, the MFGP data fusion was employed. Leveraging the high-fidelity data from the hub probe, the MFGP model extended the flow field assessment to unmeasured regions, providing a complete domain flow assessment. This approach, although not flawless, demonstrated the potential to reduce the instrumentation effort and testing times.

In the subsequent adaptive sampling simulation of the full domain, the evaluated metrics show a continuous decrease with an increasing number of measurements. A transition occurred around 230 measurements, signifying a shift towards exploration as the model concentrated on reducing uncertainty rather than improving the mean prediction.

Visualizing the simulated pressure field at various stages of measurement acquisition emphasized the methodology effectiveness. Around 230 measurements, critical flow

features were correctly identified, and subsequent measurements primarily targeted high uncertainty areas.

In summary, the adaptive *in situ* sampling strategy showcased its ability to adapt to complex experimental setups, recognize essential flow structures, and dynamically adjust sampling efforts. The integration of multi-fidelity modelling addressed data limitations, presenting a promising avenue for exploring regions of the flow that cannot be sampled. The simulation results provided insights into the adaptive methodology behaviour and its performance, showcasing its adaptability and efficiency in characterizing complex flow phenomena in turbomachinery applications.

Chapter 6.

Conclusions and Future Perspectives

In this work, an experimental measurement technique was developed and validated, systematically reducing both the instrumentation effort and testing time necessary for comprehensive flow assessments.

The methodology, couples standard experimental measurements and numerical simulations within a Bayesian data fusion framework driven by Multi-Fidelity Gaussian Processes. A novel DoE approach based on Sparse Variational Gaussian Processes and CFD simulations, was seamlessly integrated into the early phases of an experimental campaign design, guiding how much and where to measure.

The methodology results in a substantial reduction of at least 50% in measurement points and a relevant one-hour cutback per test, without loss of information, hence reshaping experimental campaign planning. This reduction not only marks an advancement in experimental sampling efficiency but also directly translates into substantial energy and personnel cost savings. The correlation between reduced testing time and lower energy consumption underscores the practical and economic impact of this methodology.

The introduction of *In situ* adaptive sampling adds a layer of adaptability to the methodology. This dynamic enhancement to measurement accuracy during experimental sampling reinforces the methodology's versatility, enabling it to tackle unforeseen challenges that might arise during an experimental campaign.

The modular nature of both sampling approaches adds another dimension of flexibility, making the methodology adaptable to different test cases. This ensures that the methodology can achieve the highest accuracy within time or measurement constraints.

Despite the varying levels of flow complexity and acquisition challenges presented by the H25 compressor and the ECL5 fan, that impact the estimated uncertainty, the methodology's exhibits generalizability. Its applicability across industrial and academic settings positions it as a valuable asset for researchers, engineers, and industries alike,

representing a significant advancement in experimental methodologies for aerodynamic assessments of turbomachinery components.

Future Perspectives

While the focus of this work has been on establishing a robust forward testing procedure, it is important to recognize that the methodology and tools developed offer opportunities for further exploitation and improvement.

Multi-fidelity exploitation for boundary layer inference

The successful application of the MFGP framework has laid a robust foundation for future advancements. One promising approach for further exploration involves extending the framework capabilities to infer flow quantities in areas traditionally challenging for instrumentation, such as boundary layers. This expansion not only addresses the limitations of current measurement techniques but also opens new possibilities for gaining deeper insights into complex flow phenomena.

As an example, from the H25 axial compressor experimental database, the last radial measurements, corresponding to the probe mechanical limit, are removed from the database both at hub and casing. A MFGP model, using all the other available measurements is trained. The model inference of the pressure tangential profile at the removed hub ($r = 0.02$) and casing ($r = 0.98$) is presented below in Figure 6.1.

This preliminary results presents a match between MFGP inference and the experiments, not used in the model training. This shows the capability of extrapolation of the GP model supported by a multi fidelity modelling framework and opens the possibility of exploiting this framework to infer flow quantities inside the boundary layer without having to resort to extra instrumentation that is usually custom made and delicate to operate with accuracy.

This approach provides new challenges, since it increases the number of variables being modelled simultaneously, for example to infer a velocity profile, but also may require the need to incorporate different instrumentations in one model with different associated noises. However, the successful extrapolation showcased in the H25 axial compressor example show promise.

Extending the MFGP framework to infer flow quantities in traditionally challenging areas, like boundary layers, offers an opportunity beyond current measurements and it has the potential to reduce instrumentation intrusiveness in such an important, but sensitive flow region that it is hard to measure accurately in an industrial rig.




Figure 6.1. MFGP tangential pressure profile inference at the H25 compressor probe sampled limits.

Performance evaluation and Rake measurements

A final proposition that this work leaves open is the optimal evaluation of the stage performance. Leveraging correct mean flow inference enhances spatial averaging evaluation, presenting opportunities for improved performance evaluation, particularly in comparison with rake measurements.

This work leaves open the leveraging of this information to improve the evaluation of the performance by the rakes. Figure 6.2 shows a preliminary possibility of exploiting the proposed methodology tools to improve the evaluation of the complete operating fan map. The Figure shows lines of increasing rotational speed, from 55NM to 105NM.

A MFGP model is applied with mass-averaged rake measurements serving as high-fidelity data and mass-averaged CFD simulations as low-fidelity data. Rake measurements from 55NM to 100NM (shown in red) are used for training the MFGP model, while measurements for 105NM (shown in green) are reserved for validation.

The inferred operating curves align well with the rake measurements, even for the 105NM rotational speed, for which no data was used in model training. This suggests that the MFGP model effectively leverages available data to learn the relationship between CFD simulations and experimental measurements, thereby extrapolating the operating curve accurately, within some level of uncertainty. This indicates potential approaches for refining rake measurements and reducing associated uncertainty, as well




Figure 6.2. MFGP model compressor map using mass-averaged CFD simulations and mass-averaged rake measurements.

as extracting otherwise inaccessible information.

Furthermore, in future work, the proposed methodology could be adapted to determine optimal positions for performance-measuring rakes. Assuming a full annulus simulation is available, the sampling methods developed here could be used to infer what are the optimal rake positions to obtain the most accurate performance measurements using these type of instrumentation. This could be performed by 'simulating' an adaptive sampling rake acquisition test, in the full annulus CFD, where the stopping criteria is set on the convergence of the estimated average pressure.

Optimal instrumentation displacement

In the application of the methodology to the ECL5 fan, a one-hour time reduction was achieved compared to the reference, fully automated test, which spanned over three hours. This significant improvement, however, was obtained with a 50% reduction in instrumentation.

To further expedite testing processes, a strategic exploration into the optimal displacement of instrumentation within the flow field is imperative. The SVGP DoE involves sampling a sparse grid, introducing the opportunity to plan an optimal acqui-

sition path.

Essentially, this optimization task aligns with the classical "travelling salesman problem". Given the full sparse grid and the distances between sampling points, the objective is to compute the shortest possible acquisition route. This optimization promises a more substantial reduction in testing time, thus directly contributing to enhancing the efficiency of the methodology in an experimental campaign.

Adaptive sampling batch acquisition

During the ECL5 experimental campaign, to take advantage of the dual parallel probes available, the acquisition function was slightly altered to acquire two values at each iteration. The proposed approach here was not optimal and for the future work an optimal acquisition strategy that can leverage all the instrumentation available in an optimal manner should be developed.

In the context of Bayesian optimization, for which the proposed adaptive sampling methodology was derived, the acquisition of multiple design points, or samples in the context of this work, is called batch sampling. ZHAN et al., 2017, in the context of optimization problems, proposes the introduction of an 'influence function' (IF) to select multiple sample points for each acquisition iteration. Their proposed IF is:

$$IF(\mathbf{x}, \mathbf{x}^{q-1}) = \prod_{i=1}^{q-1} [1 - \mathbf{K}(\mathbf{x}, \mathbf{x}_i)] \quad (6.1)$$

where \mathbf{x}_i is the i optimal point (measurement location), and $\mathbf{K}(\mathbf{x}, \mathbf{x}_i)$ is the covariance function of the initial optimal point with the rest of the candidate domain points. The IF then tends to benefit the sampling of points far from the initial optimal point.

Another option would be to rely on the hyperparameters trained by the GP model. The lengthscale of the kernel relates the amount of information that is propagated between domain points. If this value can be leveraged, an influence function could be developed based on points that have little information associated with them.

The introduction of a new step in the iteration loop, while computationally feasible, necessitates considerations regarding instrumentation translational issues and potential delays in the acquisition process. Moreover, the trade off between acquiring more measurements per iteration but these locations being suboptimal could hinder the acquisition process more expensive.

The future perspectives outlined here collectively contribute to improving the proposed measurement approach and advancing turbomachinery experimental research methodologies.

The approach presented in this work shows that, while the recent machine learning data driven trends still pose challenges, they also present opportunities for new research paths.

Bibliography

References are sorted in alphabetical order.

- ASCH, M., M. BOCQUET, and M. NODET (2016). *Data Assimilation: Methods, Algorithms, and Applications*. Volume 28. PB, pages 123–155. ISBN: 9781611974539 (cited on page 12).
- BABAEI, H., P. PERDIKARIS, C. CHRYSOSTOMIDIS, and G. E. KARNIADAKIS (2016). “Multi-fidelity modelling of mixed convection based on experimental correlations and numerical simulations”. In: *Journal of Fluid Mechanics* 809, pages 895–917. ISSN: 14697645. DOI: [10.1017/jfm.2016.718](https://doi.org/10.1017/jfm.2016.718) (cited on page 31).
- BABIN, C. (2022). “Impact of shrouded stator cavity flow on axial compressor performance and stability”. Theses. Université de Lyon (cited on pages 62, 65).
- BALLAL, D. R. and J. ZELINA (2004). “Progress in Aeroengine Technology (1939-2003)”. In: *Journal of Aircraft* 41.1, pages 43–50. ISSN: 15333868. DOI: [10.2514/1.562](https://doi.org/10.2514/1.562) (cited on page 2).
- BANNISTER, R. N. (2017). *A review of operational methods of variational and ensemble-variational data assimilation*. DOI: [10.1002/qj.2982](https://doi.org/10.1002/qj.2982) (cited on page 14).
- BERTRAM, A., P. BEKEMEYER, and M. HELD (2021). “Fusing distributed aerodynamic data using Bayesian Gappy Proper Orthogonal Decomposition”. In: *AIAA Aviation and Aeronautics Forum and Exposition, AIAA AVIATION Forum 2021*. Reston, Virginia: American Institute of Aeronautics and Astronautics. ISBN: 9781624106101. DOI: [10.2514/6.2021-2602](https://doi.org/10.2514/6.2021-2602) (cited on page 8).
- BISHOP, C. (2006). “Pattern recognition and machine learning”. In: *Springer* 2, pages 35–42 (cited on page 22).
- BLUNDELL, C., J. CORNEBISE, K. KAVUKCUOGLU, and D. WIERSTRA (2015). “Weight Uncertainty in Neural Networks”. In: *32nd International Conference on Machine Learning, ICML 2015* 2, pages 1613–1622. arXiv: [1505.05424](https://arxiv.org/abs/1505.05424) (cited on page 22).
- BRANDSTETTER, C., V. PAGES, P. DUQUESNE, B. PAOLETTI, S. AUBERT, and X. OTTAVY (2019). “Project PHARE-2—A High-Speed UHBR Fan Test Facility for a

- New Open-Test Case”. In: *Journal of Turbomachinery* 141.10. ISSN: 15288900. DOI: [10.1115/1.4043883](https://doi.org/10.1115/1.4043883) (cited on page 101).
- BREVAULT, L., M. BALESSENT, and A. HEBBAL (2020). “Overview of Gaussian process based multi-fidelity techniques with variable relationship between fidelities, application to aerospace systems”. In: *Aerospace Science and Technology* 107, page 106339. ISSN: 12709638. DOI: [10.1016/j.ast.2020.106339](https://doi.org/10.1016/j.ast.2020.106339). arXiv: [2006.16728](https://arxiv.org/abs/2006.16728) (cited on pages 28, 29).
- BRUNTON, S. L., B. R. NOACK, and P. KOUMOUTSAKOS (2020). “Machine Learning for Fluid Mechanics”. In: *Annual Review of Fluid Mechanics* 52, pages 477–508. ISSN: 00664189. DOI: [10.1146/annurev-fluid-010719-060214](https://doi.org/10.1146/annurev-fluid-010719-060214). arXiv: [1905.11075](https://arxiv.org/abs/1905.11075) (cited on page 19).
- BYRD, R. H., P. LU, J. NOCEDAL, and C. ZHU (1995). “A Limited Memory Algorithm for Bound Constrained Optimization”. In: *SIAM Journal on Scientific Computing* 16.5, pages 1190–1208. ISSN: 1064-8275. DOI: [10.1137/0916069](https://doi.org/10.1137/0916069) (cited on page 43).
- CAI, S., Z. WANG, F. FUEST, Y. J. JEON, C. GRAY, and G. E. KARNIADAKIS (2021). “Flow over an espresso cup: Inferring 3-D velocity and pressure fields from tomographic background oriented Schlieren via physics-informed neural networks”. In: *Journal of Fluid Mechanics* 915, pages 1–17. ISSN: 14697645. DOI: [10.1017/jfm.2021.135](https://doi.org/10.1017/jfm.2021.135). arXiv: [2103.02807](https://arxiv.org/abs/2103.02807) (cited on pages 19, 20).
- CAMBIER, L., S. HEIB, and S. PLOT (2013). “The Onera elsA CFD software: input from research and feedback from industry”. In: *Mechanics & Industry* 14.3, pages 159–174. ISSN: 2257-7777. DOI: [10.1051/meca/2013056](https://doi.org/10.1051/meca/2013056) (cited on page 108).
- CELIK, I. B., U. GHIA, P. J. ROACHE, C. J. FREITAS, H. COLEMAN, and P. E. RAAD (2008). “Procedure for estimation and reporting of uncertainty due to discretization in CFD applications”. In: *Journal of Fluids Engineering, Transactions of the ASME* 130.7, pages 0780011–0780014. ISSN: 00982202. DOI: [10.1115/1.2960953](https://doi.org/10.1115/1.2960953) (cited on page 67).
- CHILLA, M., G. PULLAN, and S. GALLIMORE (2020). “Reducing Instrumentation Errors Caused by Circumferential Flow-Field Variations in Multistage Axial Compressors”. In: *Journal of Turbomachinery* 142.9. ISSN: 0889-504X. DOI: [10.1115/1.4047610](https://doi.org/10.1115/1.4047610) (cited on page 9).
- CRUZ, G. G., C. BABIN, X. OTTAVY, and F. FONTANETO (2021). “Bayesian Inference of Experimental Data for Axial Compressor Performance Assessment”. In: *Volume 2A: Turbomachinery — Axial Flow Fan and Compressor Aerodynamics*. Volume 2A-2021. American Society of Mechanical Engineers. ISBN: 978-0-7918-8490-4. DOI: [10.1115/GT2021-59200](https://doi.org/10.1115/GT2021-59200) (cited on pages 15, 16, 17).

-
- DAMIANOU, A. C. and N. D. LAWRENCE (2013). “Deep Gaussian Processes, Damianou & Lawrence”. In: *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics (AISTATS)* 31.3, pages 514–521. ISSN: 1938-7228 (cited on page 29).
- ECMWF (2017). *20 years of 4D-Var: better forecasts through a better use of observations*. [\url{https://www.ecmwf.int/en/about/media-centre/news/2017/20-years-4d-var-better-forecasts-through-better-use-observations}](https://www.ecmwf.int/en/about/media-centre/news/2017/20-years-4d-var-better-forecasts-through-better-use-observations) (cited on page 14).
- EUROSTAT (2023). *Electricity price statistics* (cited on page 4).
- EVENSEN, G. (2010). *Data assimilation: The ensemble kalman filter*. ISBN: 3540383000. DOI: 10.1007/978-3-540-38301-7 (cited on page 15).
- FERNÁNDEZ-GODINO, M. G. (2016). “Review of multi-fidelity models”. In: *AIAA Journal* 57.5, pages 2039–2054. ISSN: 0001-1452. DOI: 10.2514/1.J057750. arXiv: 1609.07196 (cited on page 28).
- FORRESTER, A. I. J., A. SÓBESTER, and A. J. KEANE (2008). *Engineering Design via Surrogate Modelling*. DOI: 10.1002/9780470770801 (cited on page 44).
- FOURES, D. P., N. DOVETTA, D. SIPP, and P. J. SCHMID (2014). “A data-assimilation method for reynolds-averaged navier-stokes-driven mean flow reconstruction”. In: *Journal of Fluid Mechanics* 759, pages 404–431. ISSN: 14697645. DOI: 10.1017/jfm.2014.566 (cited on page 14).
- FUHG, J. N., A. FAU, and U. NACKENHORST (2021). *State-of-the-Art and Comparative Review of Adaptive Sampling Methods for Kriging*. Volume 28. 4. Springer Netherlands, pages 2689–2747. ISBN: 1183102009474. DOI: 10.1007/s11831-020-09474-6 (cited on page 53).
- GARUD, S. S., I. A. KARIMI, and M. KRAFT (2017). “Design of computer experiments: A review”. In: *Computers and Chemical Engineering* 106, pages 71–95. ISSN: 00981354. DOI: 10.1016/j.compchemeng.2017.05.010 (cited on page 52).
- GPY (2012). *{GPy}: A Gaussian process framework in python*. [\url{http://github.com/SheffieldML/GPy}](http://github.com/SheffieldML/GPy) (cited on pages 44, 49).
- LE-GRATIET, L. (2013). *Multi-fidelity Gaussian process regression for computer experiments, Ph.D. thesis*. Technical report. France, page 307 (cited on page 31).
- HAYASE, T. (2015). “Numerical simulation of real-world flows”. In: *Fluid Dynamics Research* 47.5, page 51201. ISSN: 01695983. DOI: 10.1088/0169-5983/47/5/051201 (cited on page 12).
- HENSMAN, J., N. FUSI, and N. D. LAWRENCE (2013). “Gaussian processes for big data”. In: *arXiv preprint arXiv:1309.6835* (cited on pages 32, 49).

- IGLESIAS, M. A., K. J. LAW, and A. M. STUART (2013). “Ensemble kalman methods for inverse problems”. In: *Inverse Problems* 29.4. ISSN: 02665611. DOI: [10.1088/0266-5611/29/4/045001](https://doi.org/10.1088/0266-5611/29/4/045001) (cited on page 15).
- JIN, R., W. CHEN, and A. SUDJIANTO (2005). “An efficient algorithm for constructing optimal design of computer experiments”. In: *Journal of Statistical Planning and Inference* 134.1, pages 268–287. ISSN: 03783758. DOI: [10.1016/j.jspi.2004.02.014](https://doi.org/10.1016/j.jspi.2004.02.014) (cited on pages 52, 85).
- JOSHI, S. and S. BOYD (2009). “Sensor selection via convex optimization”. In: *IEEE Transactions on Signal Processing* 57.2, pages 451–462. ISSN: 1053587X. DOI: [10.1109/TSP.2008.2007095](https://doi.org/10.1109/TSP.2008.2007095) (cited on page 47).
- KARNIADAKIS, G. E., I. G. KEVREKIDIS, L. LU, P. PERDIKARIS, S. WANG, and L. YANG (2021). “Physics-informed machine learning”. In: *Nature Reviews Physics* 3.6, pages 422–440. ISSN: 2522-5820. DOI: [10.1038/s42254-021-00314-5](https://doi.org/10.1038/s42254-021-00314-5) (cited on pages 19, 20).
- KENNEDY, M. and A. O’HAGAN (2000). “Predicting the output from a complex computer code when fast approximations are available”. In: *Biometrika* 87.1, pages 1–13. ISSN: 0006-3444. DOI: [10.1093/biomet/87.1.1](https://doi.org/10.1093/biomet/87.1.1) (cited on pages 29, 31).
- KIANIFAR, M. R. and F. CAMPEAN (2020). “Performance evaluation of metamodeling methods for engineering problems: towards a practitioner guide”. In: *Structural and Multidisciplinary Optimization* 61.1, pages 159–186. ISSN: 16151488. DOI: [10.1007/s00158-019-02352-1](https://doi.org/10.1007/s00158-019-02352-1) (cited on page 40).
- KRAUSE, A., A. SINGH, and C. GUESTRIN (2008). “Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies”. In: *Journal of Machine Learning Research* 9, pages 235–284. ISSN: 15324435. DOI: [10.1145/1390681.1390689](https://doi.org/10.1145/1390681.1390689) (cited on page 47).
- LAW, K. J. and A. M. STUART (2012). “Evaluating Data Assimilation Algorithms”. In: *Monthly Weather Review* 140.11, pages 3757–3782. ISSN: 00270644. DOI: [10.1175/MWR-D-11-00257.1](https://doi.org/10.1175/MWR-D-11-00257.1). arXiv: [1107.4118](https://arxiv.org/abs/1107.4118) (cited on page 14).
- LE GRATIET, L. and J. GARNIER (2014). “RECURSIVE CO-KRIGING MODEL FOR DESIGN OF COMPUTER EXPERIMENTS WITH MULTIPLE LEVELS OF FIDELITY”. In: *International Journal for Uncertainty Quantification* 4.5, pages 365–386. ISSN: 2152-5080. DOI: [10.1615/Int.J.UncertaintyQuantification.2014006914](https://doi.org/10.1615/Int.J.UncertaintyQuantification.2014006914) (cited on page 31).
- LIU, H., J. CAI, and Y. S. ONG (2017). “An adaptive sampling approach for Kriging metamodeling by maximizing expected prediction error”. In: *Computers and Chemical*

-
- Engineering* 106, pages 171–182. ISSN: 00981354. DOI: [10.1016/j.compchemeng.2017.05.025](https://doi.org/10.1016/j.compchemeng.2017.05.025) (cited on page 53).
- LOU, F. and N. L. KEY (2021). “Reconstructing Compressor Non-Uniform Circumferential Flow Field From Spatially Undersampled Data—Part 1: Methodology and Sensitivity Analysis”. In: *Journal of Turbomachinery* 143.8, pages 1–12. ISSN: 0889-504X. DOI: [10.1115/1.4050433](https://doi.org/10.1115/1.4050433) (cited on pages 8, 9, 11).
- LOU, F., D. R. MATTHEWS, N. J. KORMANIK, and N. L. KEY (2021). “Reconstructing Compressor Nonuniform Circumferential Flow Field From Spatially Undersampled Data—Part 2: Practical Application for Experiments”. In: *Journal of Turbomachinery* 143.8, pages 1–12. ISSN: 0889-504X. DOI: [10.1115/1.4050434](https://doi.org/10.1115/1.4050434) (cited on page 8).
- MANOHAR, K., B. W. BRUNTON, J. N. KUTZ, and S. L. BRUNTON (2018). “Data-Driven Sparse Sensor Placement for Reconstruction: Demonstrating the Benefits of Exploiting Known Patterns”. In: *IEEE Control Systems* 38.3, pages 63–86. ISSN: 1066033X. DOI: [10.1109/MCS.2018.2810460](https://doi.org/10.1109/MCS.2018.2810460) (cited on page 47).
- MENTER, F. R. (1994). “Two-equation eddy-viscosity turbulence models for engineering applications”. In: *AIAA Journal*. ISSN: 00011452. DOI: [10.2514/3.12149](https://doi.org/10.2514/3.12149) (cited on page 67).
- OKADA, M., L. SIMONASSI, G. LOPES, and S. LAVAGNOLI (2024). “Particle Image Velocimetry Measurements in a High-Speed Low-Reynolds Low-Pressure Turbine Cascade”. In: *Journal of Turbomachinery* 146.3, page 31010. ISSN: 0889-504X. DOI: [10.1115/1.4063674](https://doi.org/10.1115/1.4063674) (cited on page 8).
- PALEYES, A., M. PULLIN, M. MAHSERECI, N. LAWRENCE, and J. GONZÁLEZ (2019). “Emulation of physical processes with Emukit”. In: *Second Workshop on Machine Learning and the Physical Sciences, NeurIPS* (cited on page 44).
- PEHERSTORFER, B., K. WILLCOX, and M. GUNZBURGER (2018). “Survey of multifidelity methods in uncertainty propagation, inference, and optimization”. In: *SIAM Review* 60.3, pages 550–591. ISSN: 00361445. DOI: [10.1137/16M1082469](https://doi.org/10.1137/16M1082469). arXiv: [1806.10761](https://arxiv.org/abs/1806.10761) (cited on page 28).
- PERDIKARIS, P., M. RAISSI, A. DAMIANOU, N. D. LAWRENCE, and G. E. KARNIADAKIS (2017). “Nonlinear information fusion algorithms for data-efficient multifidelity modelling”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 473.2198. ISSN: 14712946. DOI: [10.1098/rspa.2016.0751](https://doi.org/10.1098/rspa.2016.0751) (cited on page 29).
- PINO, F., L. SCHENA, J. RABAULT, and M. A. MENDEZ (2023). “Comparative analysis of machine learning methods for active flow control”. In: *Journal of Fluid Mechanics* 958, A39. ISSN: 0022-1120. DOI: [10.1017/jfm.2023.76](https://doi.org/10.1017/jfm.2023.76) (cited on page 19).

- RAIOLA, M., S. DISCETTI, and A. IANIRO (2015). “On PIV random error minimization with optimal POD-based low-order reconstruction”. In: *Experiments in Fluids* 56.4, page 75. ISSN: 0723-4864. DOI: [10.1007/s00348-015-1940-8](https://doi.org/10.1007/s00348-015-1940-8) (cited on page 8).
- RASMUSSEN, C. E. and C. K. I. WILLIAMS (2005). *Gaussian Processes for Machine Learning*. Volume 7. 5. The MIT Press, pages 32–46. ISBN: 9780262256834. DOI: [10.7551/mitpress/3206.001.0001](https://doi.org/10.7551/mitpress/3206.001.0001) (cited on page 23).
- SACKS, J., W. J. WELCH, T. J. MITCHELL, and H. P. WYNN (1989). “Design and Analysis of Computer Experiments”. In: *Statistical Science* 4.4, pages 45–52. ISSN: 0883-4237. DOI: [10.1214/ss/1177012413](https://doi.org/10.1214/ss/1177012413) (cited on page 52).
- SCHMELZER, M., R. P. DWIGHT, and P. CINNELLA (2020). “Discovery of Algebraic Reynolds-Stress Models Using Sparse Symbolic Regression”. In: *Flow, Turbulence and Combustion* 104.2-3, pages 579–603. ISSN: 1386-6184. DOI: [10.1007/s10494-019-00089-x](https://doi.org/10.1007/s10494-019-00089-x). arXiv: [1905.07510](https://arxiv.org/abs/1905.07510) (cited on page 12).
- SCHNEIDER, A. P., A.-L. FIQUET, B. PAOLETTI, X. OTTAVY, and C. BRANDSTETTER (2023). “Experiments on Tuned UHBR Open-Test-Case Fan ECL5/CATANA: Performance and Aerodynamics”. In: *Volume 13A: Turbomachinery — Axial Flow Fan and Compressor Aerodynamics*. American Society of Mechanical Engineers, pages 1–12. ISBN: 978-0-7918-8708-0. DOI: [10.1115/GT2023-103629](https://doi.org/10.1115/GT2023-103629) (cited on pages 101, 103).
- SEBASTIANI, P. and H. P. WYNN (2000). “Maximum entropy sampling and optimal Bayesian experimental design”. In: *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 62.1, pages 145–157. ISSN: 13697412. DOI: [10.1111/1467-9868.00225](https://doi.org/10.1111/1467-9868.00225) (cited on page 47).
- SESHADRI, P., A. DUNCAN, and G. THORNE (2022). “Bayesian Mass Averaging in Rigs and Engines”. In: *Journal of Turbomachinery* 144.8. ISSN: 0889-504X. DOI: [10.1115/1.4053534](https://doi.org/10.1115/1.4053534). arXiv: [2011.09240](https://arxiv.org/abs/2011.09240) (cited on pages 10, 11, 46).
- SESHADRI, P., A. B. DUNCAN, G. THORNE, G. PARKS, R. VAZQUEZ DIAZ, and M. GIROLAMI (2022). “Bayesian assessments of aeroengine performance with transfer learning”. In: *Data-Centric Engineering* 3. DOI: [10.1017/dce.2022.29](https://doi.org/10.1017/dce.2022.29). arXiv: [2011.14698](https://arxiv.org/abs/2011.14698) (cited on pages 10, 11).
- SINGH, A. P., K. DURAISAMY, and Z. J. ZHANG (2017). “Augmentation of Turbulence Models Using Field Inversion and Machine Learning”. In: *55th AIAA Aerospace Sciences Meeting*. January. Reston, Virginia: American Institute of Aeronautics and Astronautics, pages 1–18. ISBN: 978-1-62410-447-3. DOI: [10.2514/6.2017-0993](https://doi.org/10.2514/6.2017-0993) (cited on page 12).

-
- SOUSA, J., C. GARCÍA-SÁNCHEZ, and C. GORLÉ (2018). “Improving urban flow predictions through data assimilation”. In: *Building and Environment* 132. January, pages 282–290. ISSN: 03601323. DOI: [10.1016/j.buildenv.2018.01.032](https://doi.org/10.1016/j.buildenv.2018.01.032) (cited on pages 15, 16, 46).
- SOUSA, J. and C. GORLÉ (2019). “Computational urban flow predictions with Bayesian inference: Validation with field data”. In: *Building and Environment* 154. November 2018, pages 13–22. ISSN: 03601323. DOI: [10.1016/j.buildenv.2019.02.028](https://doi.org/10.1016/j.buildenv.2019.02.028) (cited on page 15).
- SUNDARARAJAN, S. and S. S. KEERTHI (2001). “Predictive Approaches for Choosing Hyperparameters in Gaussian Processes”. In: *Neural Computation* 13.5, pages 1103–1118. ISSN: 0899-7667. DOI: [10.1162/08997660151134343](https://doi.org/10.1162/08997660151134343) (cited on page 54).
- SYMON, S., N. DOVETTA, B. J. MCKEON, D. SIPP, and P. J. SCHMID (2017). “Data assimilation of mean velocity from 2D PIV measurements of flow over an idealized airfoil”. In: *Experiments in Fluids* 58.5. ISSN: 07234864. DOI: [10.1007/s00348-017-2336-8](https://doi.org/10.1007/s00348-017-2336-8) (cited on page 15).
- TAJNAFOI, G., R. ARCUCCI, L. MOTTET, C. VOURIOT, M. MOLINA-SOLANA, C. PAIN, and Y. K. GUO (2021). “Variational Gaussian Process for Optimal Sensor Placement”. In: *Applications of Mathematics* 66.2, pages 287–317. ISSN: 15729109. DOI: [10.21136/AM.2021.0307-19](https://doi.org/10.21136/AM.2021.0307-19) (cited on page 47).
- TAYLOR, J. V. and R. J. MILLER (2017). “Competing three-dimensional mechanisms in compressor flows”. In: *Journal of Turbomachinery* 139.2, pages 1–10. ISSN: 15288900. DOI: [10.1115/1.4034685](https://doi.org/10.1115/1.4034685) (cited on page 68).
- VENTURI, D. and G. E. KARNIADAKIS (2004). “Gappy data and reconstruction procedures for flow past a cylinder”. In: *Journal of Fluid Mechanics* 519, pages 315–336. ISSN: 00221120. DOI: [10.1017/S0022112004001338](https://doi.org/10.1017/S0022112004001338) (cited on pages 8, 11).
- VETRA-CARVALHO, S., P. J. VAN LEEUWEN, L. NERGER, A. BARTH, M. U. ALTAF, P. BRASSEUR, P. KIRCHGESSNER, and J. M. BECKERS (2018). “State-of-the-art stochastic data assimilation methods for high-dimensional non-Gaussian problems”. In: *Tellus, Series A: Dynamic Meteorology and Oceanography* 70.1, pages 1–38. ISSN: 16000870. DOI: [10.1080/16000870.2018.1445364](https://doi.org/10.1080/16000870.2018.1445364) (cited on page 14).
- VIRTANEN, P. et al. (2020). “SciPy 1.0: fundamental algorithms for scientific computing in Python”. In: *Nature Methods* 17.3, pages 261–272. ISSN: 1548-7091. DOI: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2). arXiv: [1907.10121](https://arxiv.org/abs/1907.10121) (cited on page 43).
- WIKLE, C. K. and L. M. BERLINER (2007). “A Bayesian tutorial for data assimilation”. In: *Physica D: Nonlinear Phenomena* 230.1-2, pages 1–16. ISSN: 01672789. DOI: [10.1016/j.physd.2006.09.017](https://doi.org/10.1016/j.physd.2006.09.017) (cited on page 21).

- YU, X. J. and B. J. LIU (2007). “Stereoscopic PIV measurement of unsteady flows in an axial compressor stage”. In: *Experimental Thermal and Fluid Science* 31.8, pages 1049–1060. ISSN: 08941777. DOI: [10.1016/j.expthermflusci.2006.11.001](https://doi.org/10.1016/j.expthermflusci.2006.11.001) (cited on page 8).
- ZAMBONINI, G., X. OTTAVY, and J. KRIEGSEIS (2017). “Corner Separation Dynamics in a Linear Compressor Cascade”. In: *Journal of Fluids Engineering* 139.6. ISSN: 0098-2202. DOI: [10.1115/1.4035876](https://doi.org/10.1115/1.4035876) (cited on page 8).
- ZHAN, D., J. QIAN, and Y. CHENG (2017). “Pseudo expected improvement criterion for parallel EGO algorithm”. In: *Journal of Global Optimization* 68.3, pages 641–662. ISSN: 0925-5001. DOI: [10.1007/s10898-016-0484-7](https://doi.org/10.1007/s10898-016-0484-7) (cited on page 145).
- ZHANG, X. L., C. MICHELÉN-STRÖFER, and H. XIAO (2020). “Regularized ensemble Kalman methods for inverse problems”. In: *Journal of Computational Physics* 416, page 109517. ISSN: 10902716. DOI: [10.1016/j.jcp.2020.109517](https://doi.org/10.1016/j.jcp.2020.109517). arXiv: [1910.01292](https://arxiv.org/abs/1910.01292) (cited on page 22).