



UNIVERSITY OF PAVIA
FACULTY OF ENGINEERING
DEPARTMENT OF CIVIL ENGINEERING AND ARCHITECTURE

DOCTOR OF PHILOSOPHY
XXXVIII CYCLE

Uncertainty-Aware Surrogate Modeling and Bayesian Inverse Design for Process-Structure-Property Optimization

Applications to Metal Additive Manufacturing, Architected Materials, and Structure Optimization

Supervisors:

Prof. Ferdinando Auricchio
University of Pavia
Prof. Massimo Carraturo
University of Pavia

Candidate:

Mihaela Chiappetta, M.Sc.
MAT. 526809

A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy
in
Design, Modeling and Simulation in Engineering

Abstract

Process–structure–property relationships govern predictive modeling and performance-driven design across a wide range of engineering applications, including metal additive manufacturing, architected materials, and structural optimization. Their rigorous characterization remains challenging due to the multiscale nature of the underlying physical processes, variability in material properties and fabrication conditions, limited availability of experimental and numerical data, and the presence of both epistemic and aleatory uncertainties.

In metal additive manufacturing, strong thermal gradients and melt-pool dynamics induce anisotropic mechanical responses, microstructural heterogeneity, and residual strains, leading to residual stresses, geometric distortions, and defects in manufactured components. Such effects reduce process repeatability and hinder both experimental characterization and predictive modeling.

Architected materials such as spinodoids exhibit high-dimensional design spaces governed by nonlinear and highly sensitive structure–property mappings. Small variations in geometric parameters or fabrication tolerances may induce significant changes in effective mechanical properties, thereby limiting systematic exploration and robust design.

Structural optimization introduces additional challenges, as governing partial differential equations depend on evolving design variables and must remain consistent with nonlinear material behavior and manufacturing constraints. At the same time, high-fidelity simulations and experimental investigations remain computationally and experimentally expensive, restricting the systematic exploration of the parameter space and limiting the data available for model calibration and validation.

The present doctoral thesis addresses these challenges through three complementary research directions: uncertainty quantification, inverse analysis and design, and structural optimization.

The first research direction develops an uncertainty quantification framework for thermomechanical modeling in metal additive manufacturing. The approach combines variance-based global sensitivity analysis, Bayesian inversion for parameter calibration, and forward uncertainty propagation. High-dimensional model responses are approximated through single- and multi-fidelity sparse-grid surrogate models, enabling efficient characterization of variability in process parameters and its impact on quantities of interest.

The second research direction focuses on inverse problems and inverse design through Bayesian inversion combined with surrogate modeling. Gaussian-process surrogate models are constructed via adaptive sampling driven by predictive uncertainty, enabling efficient approximation of computationally expensive forward models. The inverse problem is formulated within a Bayesian framework, where posterior distributions are approximated through local Laplace approximations

around maximum a posteriori estimates, providing both parameter identification and uncertainty quantification.

The third research direction investigates structural topology and nonlinear optimization methods, with particular emphasis on the Solid Isotropic Material Penalization (SIMP) approach combined with Optimality Criteria and on the Method of Moving Asymptotes. The analysis provides a rigorous formulation and a critical assessment of these established methodologies, clarifying their theoretical structure and numerical behavior on benchmark problems, and establishing a consistent methodological basis for future extensions toward uncertainty-aware optimization frameworks.

The proposed methodologies are applied to thermomechanical calibration in metal additive manufacturing under uncertainty, to the inverse design of spinodoid architected materials targeting prescribed effective properties, and to benchmark problems in structural topology optimization.

Although developed independently, the three research directions share common methodological components, including surrogate modeling, approximation of high-dimensional responses, and iterative solution strategies for computationally expensive problems. These common elements indicate the potential for integrating forward uncertainty quantification, Bayesian inversion, and structural optimization within a unified computational framework for uncertainty-aware design.

The development of such a framework, enabling the consistent propagation of uncertainty through forward modeling, parameter inference, and design optimization, represents a natural direction for future research.

Overall, the results contribute to reducing computational cost, enabling uncertainty-informed predictions, and supporting the identification of designs that satisfy prescribed performance requirements under variability and limited data.

Acknowledgments

I would like to express my sincere gratitude to my Ph.D. supervisor, Prof. Ferdinando Auricchio, for his guidance and support throughout these years. His experience, scientific perspective, and leadership have represented a constant reference during my doctoral studies and have contributed significantly to shaping my research development.

I am deeply grateful to my co-supervisor, Prof. Massimo Carraturo, for his continuous support, availability, and for the many insightful discussions that accompanied the development of this work. His guidance has been essential in all phases of the research, from the early formulation of ideas to their final consolidation.

I would also like to thank Dr. Chiara Piazzola, Prof. Lorenzo Tamellini, and Prof. Alessandro Reali for their support, suggestions, and supervision during the initial stages of my research activity. Their contribution has played an important role in my scientific growth and in building the foundations of this work.

I am deeply grateful to Prof. Markus Kästner for his guidance, support, and for the opportunity to carry out a research period at TU Dresden. His perspective and advice have been extremely valuable, and his support during this experience has had a meaningful impact on my work and development. I also sincerely thank Dr. Alexander Rassloff for his collaboration, for the many insightful discussions, and for his support throughout the work, which have been greatly appreciated.

I would also like to thank the research group at TU Dresden for the stimulating and collaborative research environment. In particular, I am grateful to Ulrike, Naja, Abel and Max for their support, helpful interactions, and for contributing to a welcoming and constructive atmosphere during my stay. I also acknowledge Dr. Alexandra Otto for her assistance and valuable discussions.

I would like to acknowledge the CompMech group in Pavia. Although my work has been primarily focused on individual research activities, the seminars, scientific discussions, and group initiatives have provided a stimulating academic environment and have contributed to my development during these years.

I also wish to thank Prof. Stefano Marelli and Prof. Jason Papaioannou for their careful review of my work and for their valuable comments and suggestions, which have helped improve its quality and clarity.

I am grateful to my friends for their constant presence and support. In particular, I would like to thank Silvia, Paola, Marina, and Fabiola for their encouragement, understanding, and for always being there during both the most challenging and the most rewarding moments of this period.

Finally, I would like to express my deepest gratitude to my father, who has been my constant point of reference throughout these years. His support, patience, and encouragement have sustained me at every stage of this journey. His presence has given me the strength to continue, even in the most demanding moments, and this work would not have been possible without him.

I would also like to thank my aunt Anna, Anto, and Saro for their continuous support and for always being present with care and encouragement. Their closeness has meant a great deal over time. A special acknowledgment goes to my cousin Erasmo for his valuable advice, thoughtful suggestions, and for the many conversations that have helped me reflect and move forward with greater awareness.

Mihaela Chiappetta, M.Sc.

April 24, 2026

Contents

Abstract	I
Acknowledgments	III
List of Abbreviations	XX
List of Symbols	XXII
1 Introduction	1
1.1 Motivations	1
1.2 Objectives	2
1.3 Structure and Organization of the Thesis	3
2 Uncertainty quantification of a part-scale thermomechanical cantilever beam model produced by metal additive manufacturing	4
2.1 Introduction	4
2.2 Governing equations	6
2.2.1 Thermal problem	6
2.2.2 Mechanical problem	7
2.3 Experimental setup	8
2.4 Numerical modeling	10
2.4.1 Meshing strategy	10
2.4.2 Material properties	11
2.4.3 Printing parameters	14
2.4.4 Printing process	16
2.5 Uncertainty quantification frsmework	17
2.5.1 Uncertainty quantification based on single-fidelity sparse-grid surrogate modeling	17
2.5.1.1 Sources of uncertainty	17
2.5.1.2 Uncertainty quantification workplan	19
2.5.1.3 Sparse-grid surrogate modeling for uncertainty quantification	20
2.5.1.4 Global sensitivity analysis	22
2.5.1.5 Inverse uncertainty quantification analysis	24

2.5.1.5.1	Gaussian approximation of posterior probability density function	25
2.5.1.6	Data-informed forward uncertainty quantification analysis . .	26
2.5.1.7	Results and discussion	27
2.5.1.7.1	Global sensitivity analysis	27
2.5.1.7.2	Inverse uncertainty quantification	28
2.5.1.7.2.1	Surrogate modeling	28
2.5.1.7.2.2	Bayesian inverse analysis	30
2.5.1.7.3	Data-informed forward uncertainty quantification .	34
2.5.1.7.3.1	Surrogate modeling	34
2.5.1.7.3.2	Data-informed forward analysis	35
2.5.2	Uncertainty quantification based on multi-fidelity sparse-grid surrogate modeling	37
2.5.2.1	Multi-fidelity uncertainty quantification framework	37
2.5.2.2	Multi-fidelity sparse-grid surrogate modeling	39
2.5.2.3	Results of the uncertainty quantification workflow	41
2.5.2.3.1	Inverse uncertainty quantification analysis	42
2.5.2.3.1.1	Uncertain parameters and available data . .	42
2.5.2.3.1.2	Surrogate modeling	42
2.5.2.3.1.3	Bayesian inverse analysis	45
2.5.2.3.2	Data-informed forward uncertainty quantification analysis	48
2.5.2.3.2.1	Surrogate modeling	48
2.5.2.3.2.2	Data-informed forward analysis	49
2.5.2.3.2.3	Multi-fidelity versus high-fidelity surrogate models results	52
2.6	Conclusions	53

3 Uncertainty-Aware Bayesian Inversion and Adaptive Surrogate Modeling for Inverse Design **56**

3.1	An Uncertainty-Aware Framework for Inverse Problems through Adaptive Surrogate Modeling and Bayesian Inversion	56
3.1.1	Introduction	56
3.1.2	Related Work	58
3.1.2.1	Surrogate Modeling for Inverse Problems	58
3.1.2.2	Adaptive Sampling and Uncertainty-Driven Strategies	59
3.1.2.3	Bayesian Inversion for Parameter Inference	59
3.1.3	Proposed Uncertainty-Aware Framework	60
3.1.3.1	Problem Formulation	61
3.1.3.2	Adaptive Surrogate Modeling	61
3.1.3.3	Bayesian Inversion Framework	62
3.1.4	Numerical Validation on Analytical Benchmarks	63

3.1.4.1	One-Dimensional Benchmarks	64
3.1.4.1.1	Adaptive Surrogate Model Construction	65
3.1.4.1.2	Bayesian Inversion Results	68
3.1.4.1.3	Discussion	70
3.1.4.2	Two-Dimensional Benchmarks	71
3.1.4.2.1	Adaptive Surrogate Model Construction	72
3.1.4.2.2	Bayesian Inversion Results	74
3.1.4.2.3	Discussion	79
3.1.5	Conclusions	80
3.2	Uncertainty-Aware Bayesian Inversion Using Adaptive Gaussian-Process Surrogate Modeling for the Inverse Design of Spinodoid Architected Materials	81
3.2.1	Introduction	81
3.2.2	Spinodoid Architected Materials	83
3.2.3	High-Fidelity Forward Model	85
3.2.3.1	Generation of Spinodoid Geometry	86
3.2.3.2	Numerical Homogenization	87
3.2.3.3	Extraction of Effective Mechanical Properties	87
3.2.3.4	Construction of the Structure–Property Dataset	88
3.2.4	Uncertainty-Aware Bayesian Inversion Framework	88
3.2.4.1	Adaptive Construction of Gaussian-Process Surrogate Models	89
3.2.4.2	Formulation of the Bayesian Inversion Problem	90
3.2.5	Results and Discussion	92
3.2.5.1	Results in Two-Dimensional Parameter Space	93
3.2.5.1.1	Adaptive Construction of Gaussian-Process Surrogate Models	94
3.2.5.1.2	Uncertainty-Aware Bayesian Inversion Results	97
3.2.5.2	Results in Four-Dimensional Parameter Space	101
3.2.5.2.1	Adaptive Construction of Gaussian-Process Surrogate Models	102
3.2.5.2.2	Uncertainty-Aware Bayesian Inversion Results	103
3.2.5.3	Conclusions	106
3.3	Conclusions	108
4	Structural optimization strategies with optimality criteria and method of moving asymptotes on two- and three-dimensional benchmarks	110
4.1	Introduction	110
4.2	Topology Optimization via the Solid Isotropic Material Penalization Method and an Optimality Criterion Approach	112
4.2.1	Continuum Formulation	112
4.2.1.1	Linear Elastic Problem	113
4.2.1.2	Minimum Compliance Problem	114
4.2.1.3	Lagrangian Formulation	115

4.2.1.4	Optimality Criterion	117
4.2.1.5	Regularization Techniques	118
4.2.2	Discrete Formulation	119
4.2.2.1	Linear Elastic Problem	120
4.2.2.2	Minimum Compliance Problem	121
4.2.2.3	Lagrangian Formulation	122
4.2.2.4	Optimality Criterion	123
4.2.2.5	Regularization Techniques	124
4.2.3	Numerical Implementation	125
4.2.4	Numerical Examples	127
4.2.4.1	Cantilever Beam	129
4.2.4.2	2D Short Cantilever Beam	129
4.2.4.3	L-Shaped Structure	129
4.2.5	Results and Discussion	132
4.2.5.1	Cantilever Beam	133
4.2.5.2	2D Short Cantilever Beam	133
4.2.5.3	L-Shaped Structure	134
4.3	A Comprehensive Theoretical and Algorithmic Study on the Method of Moving Asymptotes for Structural Optimization	136
4.3.1	Constrained Structural Optimization Problem: Formulation and Analytical Motivation	137
4.3.1.1	General Mathematical Formulation	138
4.3.1.2	Analytical Difficulties in Nonlinear and Nonconvex Problems	138
4.3.1.3	Analytical Motivation for Local and Separable Approximations	139
4.3.2	The Method of Moving Asymptotes: Theoretical Foundations	141
4.3.2.1	One-Dimensional Unconstrained Structural Optimization Problem	141
4.3.2.1.1	Construction of the Convex Approximation	142
4.3.2.1.2	Determination of Approximation Parameters	143
4.3.2.1.3	Adaptive Update of the Moving Asymptotes	145
4.3.2.1.4	Analytical Relation to Gradient-Based Methods	146
4.3.2.2	Extension to One-Dimensional Constrained Structural Optimization Problems	147
4.3.2.2.1	Convex Approximation of the Constraint Function	147
4.3.2.2.2	Lagrangian Formulation of the One-Dimensional Constrained Structural Optimization Problem	148
4.3.2.3	Extension to Multidimensional and Large-Scale Problems	150
4.3.2.3.1	Analytical Generalization and Separable Structure	150
4.3.3	Algorithmic Structure and Computational Implementation	152
4.3.3.1	Lagrangian Formulation and Optimality Conditions	153
4.3.3.2	Adaptive Update of the Moving Asymptotes	155
4.3.3.3	Iterative Procedure and Convergence	157

4.3.4	Numerical Validation and Analysis of Algorithmic Behaviour	157
4.3.4.1	Cantilever Beam with Displacement Constraint	158
4.3.4.2	Eight-Bar Truss under Stress Constraints	159
4.3.4.3	Two-Bar Truss with Mixed Design Variables	160
4.4	Conclusions	162
5	Conclusion and Future perspectives	167
A	Sparse-grid surrogate modeling	192
B	Leja points	195
C	Implementation Details: The UQforPy Repository	196
C.1	High-level analysis scripts	197
C.2	Algorithmic modules	198
C.3	Benchmark functions and numerical utilities	198
C.4	Repository organization and reproducibility	199
D	Adaptive Surrogate Modeling Results for One-Dimensional Benchmarks	200
E	Bayesian Inversion Results for 1D Benchmarks	202
F	Supplementary Surrogate-Based Posterior Analysis	204

List of Figures

2.1	AMBench2018-01 benchmark for the experiment performed by NIST [1, 2].	9
2.2	Part-scale model: Inconel 625 cantilever beam with length 75 mm, height 12 mm, and width 5 mm, printed on a build plate of dimensions 85 mm \times 12 mm \times 20 mm. Numbered (purple) and lettered (blue) points denote measurement locations used in the UQ analysis; the dotted (green) line at $z = 11$ mm identifies where residual strains are extracted; the solid (red) contour indicates the cut terminating at $x = 56$ mm.	11
2.3	Mesh strategies adopted for the part-scale PBF-LB/M thermomechanical model of the Inconel 625 cantilever beam. (a) Coarse mesh with characteristic element size 0.5 mm; (b) Fine mesh with characteristic element size 0.25 mm. The build plate is meshed with 3 mm elements in both cases.	12
2.4	Temperature-dependent bilinear isotropic hardening law for Inconel 625 adopted in the mechanical analysis (<i>ANSYS 2021R2</i>).	13
2.5	Temperature-dependent properties of Inconel 625 adopted in the simulations (<i>ANSYS 2021R2</i>): (a) coefficient of thermal expansion; (b) Young’s modulus; (c) Poisson’s ratio; (d) specific heat; (e) thermal conductivity.	15
2.6	Breakdown of the sparse grid (h) into the tensor grids composing it (a)–(g). Figure (i) illustrates the different families of collocation points used in this work: symmetric Leja points for a uniform PDF on $\Gamma_n = [-1, 1]$, and symmetric Gaussian Leja points for a Gaussian PDF with zero mean and unit standard deviation.	23
2.7	Magnitude of Sobol indices based on the displacements evaluated at the 11 ridges of the cantilever beam (QoI) for the three-parameter model: (a) Principal Sobol Indices that measure the <i>individual</i> contribution of each parameter to the variance of the QoI; (b) Total Sobol Indices that quantify the contributions of each parameter <i>combined with the others</i>	28
2.8	Sparse-grid surrogate model construction for the inverse UQ analysis with $w = 3$ (25 sparse-grid points) for the first ridge of the cantilever beam (i.e., $\mathbf{x}_{1,meas} = (0.5, 2.5, 12.5)$): (a) Sparse grid; (b) Surrogate model.	30
2.9	Results of the sparse-grid surrogate model for inverse UQ analysis. Convergence test for the first ridge of the cantilever beam (i.e., $\mathbf{x}_{1,meas} = (0.5, 2.5, 12.5)$): (a) Pointwise prediction error E_{PPE} ; (b) Root mean square error E_{MSE} (defined as relative error, see text).	30

2.10	Results of the sparse-grid surrogate model for inverse UQ analysis. Comparison between sparse-grid surrogate model displacements and part-scale thermomechanical model displacements for the first ridge of the cantilever beam (i.e., $\mathbf{x}_{1,meas} = (0.5, 2.5, 12.5)$): (a) Sparse grid with level $w = 1$; (b) Sparse grid with level $w = 3$	31
2.11	Displacements $u(\mathbf{x}_{k,meas}, \bar{\mathbf{v}})$ obtained from part-scale thermomechanical analysis (gray curve) and synthetic displacements data \tilde{u}_k (black marker). Error bars associated with the measurements are also reported. Unlike standard error bar plots however, the error bars are centered at the exact (yet unknown) values on the displacements $u(\mathbf{x}_{k,meas}, \bar{\mathbf{v}})$, and show the ranges within which the actual noisy measurements are most likely found, according to Equation (2.15), i.e., $u(\mathbf{x}_{k,meas}, \bar{\mathbf{v}}) \pm 3\bar{\sigma}$	32
2.12	Results of the inverse UQ analysis: (a) Surface plot of the least-squares functional $LS(\mathbf{v})$; (b) Isolines of $LS(\mathbf{v})$, target value $\bar{\mathbf{v}}$, the position of the cutting planes used to generate the one-dimensional plots in Figure 2.12(c) marked with a solid black line and in Figure 2.12(d) marked with a dashed black line, and the two MAP values computed by minimization of $LS(\mathbf{v})$; (c) LS profile at $\log h_p = -3$ with the two MAP values; (d) LS profile at $T_A = 1341^\circ\text{C}$ with the two MAP values.	33
2.13	Results of the inverse UQ analysis: (a) Uniform prior-PDF and Gaussian posterior-PDF for parameter T_A before and after Bayesian inversion; (b) Uniform PDFs for the $\log h_p$ parameter before and after Bayesian inversion.	34
2.14	Sparse-grid surrogate model construction for the data-informed forward UQ analysis with $w = 3$ (25 sparse grid points) for the first position of the cantilever beam (i.e., $\mathbf{x}_{1,str} = (1.5, 2.5, 11)$): (a) Sparse grid; (b) Surrogate model.	35
2.15	Results of the sparse-grid surrogate model for data-informed forward UQ analysis. Convergence test for the first position of the cantilever beam (i.e., $\mathbf{x}_{1,str} = (1.5, 2.5, 11)$): (a) Pointwise prediction error E_{PPE} ; (b) Root mean square error E_{MSE}	36
2.16	Results of the sparse-grid surrogate model for data-informed forward UQ analysis. Comparison between sparse-grid surrogate model residual strains and part-scale thermomechanical model residual strains for the first position of the cantilever beam (i.e., $\mathbf{x}_{1,str} = (1.5, 2.5, 11)$): (a) Sparse grid with level $w = 1$; (b) Sparse grid with level $w = 3$	36
2.17	Results of data-informed forward UQ analysis: i) the mode of the prior-based (dashed black line) and data-informed posterior (continuous red line) PDFs of the residual strain ε_{xx} ; ii) the prior-based (gray area) and data-informed posterior (pink area) uncertainty bands for the residual strain ε_{xx} ; iii) the residual strain profile (continuous black line) obtained from part-scale thermomechanical analysis for the target value $\bar{\mathbf{v}}$; The vertical dotted lines represent the 6 positions at which the prior and data-informed posterior PDFs are reported in Figure 2.18.	37

2.18	Results of data-informed forward UQ analysis. Prior-based and data-informed posterior PDFs of the residual strains at selected locations along the beam (see Figure 2.17), together with the corresponding residual strain values obtained from the part-scale thermomechanical simulation at the target value $\bar{\mathbf{v}}$	38
2.19	Residual strains obtained from the part-scale thermomechanical analysis for the target value $\bar{\mathbf{v}}$	39
2.20	PBF-LB/M numerical model of the cantilever beam 75 mm long, 12 mm high and 5 mm wide with a build platform measuring 85 mm long, 12 mm high and 20 mm wide. Points marked in purple are the locations $\mathbf{x}_{k,meas}, k = 1, \dots, 5$ used during the Bayesian inverse UQ analysis; the (blue) dotted line marks the locations where residual strains are predicted using data-informed forward UQ analysis.	42
2.21	Displacements at the 11 ridges at the top of the cantilever beam provided by NIST [1].	43
2.22	MISC surrogate model construction based on 17 evaluations (points marked in blue) of the low-fidelity PBF-LB/M numerical model ($\alpha = 1$) and 5 evaluations (points marked in yellow) of the high fidelity ($\alpha = 2$) for the first node of the cantilever beam (i.e., $\mathbf{x}_{1,meas} = (0.5, 2.5, 12.5)$): (a) sparse grid; (b) surrogate model.	44
2.23	Results of the MISC surrogate model for inverse UQ analysis. Comparison between high-fidelity displacement results with ANSYS2021-R2, U_Z , and the displacement results for three different MISC surrogate models ($\mathcal{S}_I U_Z$): a) MISC model with 6 (5+1) evaluations; b) MISC model with 10 (7+3) evaluations; c) MISC model with 22 (17+5) evaluations. The blue line represents the pairs $(U_{Z,ansys}, U_{Z,ansys})$ and the burgundy points represent the pairs $(U_{Z,ansys}, U_{Z,MISC})$, evaluated for each of the 7 validation points.	45
2.24	Results of the Bayesian inverse UQ analysis: (a) uniform prior PDF and Gaussian posterior PDF for parameter $\log h_p$ before and after Bayesian inverse UQ analysis; (b) uniform prior PDF and Gaussian posterior PDF for parameter T_A before and after Bayesian inverse UQ analysis.	47
2.25	Collocation points used for the construction of the MISC multi-fidelity surrogate models based on 5 evaluations (points marked in blue) of the PBF-LB/M numerical model at low fidelity, $\alpha = 1$, and 1 evaluation (marked in yellow) of the PBF-LB/M numerical model at high fidelity, $\alpha = 2$ for the residual strains cantilever beam model.	49
2.26	Results of the data-informed forward UQ analysis. MISC multi-fidelity surrogate models for some of the nodes of the cantilever beam, obtained from the collocation points shown in Figure 2.25.	50

2.27	Results of the data-informed forward UQ analysis. The figure includes: (i) the most likely profiles of ε_{xx} obtained from forward UQ analyses based on ρ_{prior} (the dashed grey line represents the mode of the prior-based PDFs) and ρ_{post} (the solid red line represents the mode of the posterior-based PDFs), (ii) the associated uncertainty bands, i.e., the 5%–95% quantile bands of the two PDFs (grey area for the mode prior-based PDF and pink area for the mode posterior-based PDF) and (iii) the residual strains profile provided by NIST (continuous blue line). The vertical dotted lines represent the 6 positions at which the prior and data-informed posterior PDFs are reported in Figure 2.28.	51
2.28	Results of data-informed forward UQ analysis. Prior-based and data-informed PDFs of the residual strains of the cantilever beam for the x locations marked in Figure 2.20.	52
2.29	Results of the data-informed forward UQ analysis. Top row: collocation points for the MISC surrogate model with 7 (left) and 13 (right) low-fidelity evaluations. Bottom row: corresponding modes and uncertainty bands of the residual strains.	53
2.30	Results of the data-informed forward UQ analysis. Comparison in terms of prediction of residual strains of the PBF-LB/M numerical model and associated uncertainty bands using (a) surrogate models based on multi-fidelity PBF-LB/M numerical model evaluations; (b) surrogate models based on high-fidelity PBF-LB/M numerical model evaluations.	54
3.1	Schematic representation of the proposed uncertainty-aware framework for inverse problems. The workflow consists of three stages: initialization, adaptive surrogate modeling, and Bayesian inversion. An initial dataset is used to construct a Gaussian process (GP) surrogate model, which is iteratively refined through uncertainty-aware sampling based on its predictive variance. Once a prescribed accuracy is achieved, the surrogate model is employed within a Bayesian inversion procedure, where the least-squares functional induces the likelihood and enables the estimation of the posterior distribution.	61
3.2	One-dimensional analytical benchmark functions used to validate the proposed uncertainty-aware framework for inverse problems. The suite includes representative challenges such as smooth nonlinearity, multimodal structure, and fine-scale variability, allowing systematic assessment of surrogate modeling and parameter inference capabilities. (a) Mixed Gaussian-Periodic. (b) Lévy. (c) Griewank. (d) Forrester.	65
3.3	Performance comparison of surrogate models for the Mixed Gaussian–Periodic benchmark function. (a) Predictive responses obtained with different surrogate models. The solid black line denotes the high-fidelity model; dotted lines correspond to LP (fuchsia), LE (purple), and CS (orange) deterministic alternatives; cyan and blue denote GP-based surrogate models with Matérn and RBF kernels, respectively. Red diamonds indicate number of samples. (b) MSE values for 14 number of samples. (c) MSE evolution with increasing number of samples.	67

3.4	Progressive refinement of the GP surrogate model for the Mixed Gaussian–Periodic benchmark function using the uncertainty-aware sampling. Blue solid line: high-fidelity model; green dotted line: GP mean (Matérn); light blue band: 95% confidence interval. Gray circles: initial samples; red crosses: acquired new samples. (a) 5 number of samples. (b) 10 number of samples. (c) 15 number of samples. (d) 19 number of samples.	67
3.5	Final surrogate models for all benchmark functions. Blue solid line: high-fidelity; red dotted line: GP (Matérn); light blue band: predictive uncertainty. Gray circles: initial samples; red crosses: acquired new samples. (a) Mixed Gaussian–Periodic. (b) Lévy. (c) Griewank. (d) Forrester.	68
3.6	Parameter inference via Bayesian inversion for the Mixed Gaussian–Periodic benchmark function. (a) MAP estimates obtained from multiple optimization runs. (b) LS profile across the parameter space. (c) NLS profile representing the posterior distribution approximation and highlighting the high-probability region.	69
3.7	Parameter inference via Bayesian inversion for the Forrester benchmark function. (a) MAP estimate obtained from multiple optimization runs. (b) LS profile across the parameter space. (c) NLS profile representing the posterior distribution approximation and highlighting the high-probability region.	70
3.8	Two-dimensional analytical benchmark functions used to validate the proposed uncertainty-aware framework for inverse problems. (a)–(b): Mixed Gaussian–Periodic function — high-fidelity model surface and contour plots. (c)–(d): Rosenbrock function — high-fidelity model surface and contour plots.	73
3.9	Surrogate modeling results for the Mixed Gaussian–Periodic benchmark function. (a) Final surrogate model surface after 21 high-fidelity evaluations. (b) Contour plot with sample locations. (c) Evolution of the mean squared error (MSE) as a function of the number of samples.	74
3.10	Surrogate modeling results for the Rosenbrock benchmark function. (a) Final surrogate model surface after 9 high-fidelity evaluations. (b) Contour plot with sample locations. (c) Evolution of the mean squared error (MSE) as a function of the number of samples.	75
3.11	Parameter inference via Bayesian inversion for the Mixed Gaussian–Periodic benchmark function. Red triangles: initial guesses; black square: observed quantity of interest; green stars: MAP estimates. (a) LS surface. (b) LS contour. (c) NLS surface. (d) NLS contour.	77
3.12	Parameter inference via Bayesian inversion for the Rosenbrock benchmark function. Red triangles: initial guesses; black square: observed quantity of interest; green star: MAP estimate. (a) LS surface. (b) LS contour. (c) NLS surface. (d) NLS contour.	78

3.13	Generation and parametrization of spinodoid architected materials. From left to right, the upper part illustrates the construction process: generation of the scalar Gaussian random field; restriction of wave-vector orientations via anisotropy angles θ_i ; level-set thresholding defining the solid domain and controlling the volume fraction v_f ; and rigid rotation governed by the angles ψ_i . The lower part presents representative morphologies associated with characteristic anisotropy configurations, including lamellar, columnar, cubic, and isotropic structures. The complete microstructural description is defined by the parameter vector $\Theta = (\theta_1, \theta_2, \theta_3, v_f, \psi_1, \psi_2, \psi_3)$	85
3.14	Validation NRMSE for each output component during adaptive enrichment. The error exhibits an initial transient regime followed by rapid decay and eventual stabilization, indicating progressive improvement of the surrogate approximation across all components.	95
3.15	Validation R^2 for each output component. All components rapidly approach unity, confirming the capability of the surrogate model to reproduce the high-fidelity mapping with high accuracy.	95
3.16	Aggregated validation metrics for the multi-output surrogate model. The combined indicators confirm the global accuracy and stability of the adaptive Gaussian-process approximation over the entire parameter space.	96
3.17	Gaussian-process surrogate model reconstruction of the effective stiffness components over the (θ_1, θ_2) domain.	96
3.18	Contour plots of the objective functions over the parameter space (θ_1, θ_2) . Darker regions correspond to lower objective values. The blue square denotes the reference configuration, the green star indicates the MAP estimate, and the red triangles represent the initial multistart points. The comparison highlights the effect of uncertainty-aware weighting in shaping the objective landscape while preserving the location of the global minimum.	98
3.19	Structure of the effective covariance matrix estimated from residuals using a 20-sample calibration dataset. Top: covariance and correlation matrices prior to adaptive refinement. Bottom: corresponding matrices after posterior-focused refinement.	99
3.20	Evolution of the components of the posterior-focused refinement criterion during adaptive learning (20-sample calibration case). The plots report, respectively, the predictive uncertainty at the selected points, the acquisition score, and the posterior weight driving the refinement strategy.	100
3.21	Convergence of the Gaussian-process surrogate model during adaptive enrichment in the four-dimensional parameter space. The NRMSE exhibits a rapid initial decrease followed by gradual stabilization, while R^2 quickly approaches unity and remains stable throughout the enrichment process, confirming accurate reconstruction of the high-fidelity mapping.	103

3.22	Effective covariance matrix Σ_{eff} estimated from residuals in the four-dimensional setting. Top: covariance and correlation matrices before posterior-focused refinement. Bottom: corresponding matrices after refinement.	105
3.23	Evolution of the components of the posterior-focused refinement criterion during adaptive learning.	106
4.1	2D cantilever beam with various load applications [3]. (a) Central load. (b) Load at bottom-right corner. (c) Load at top-right corner.	130
4.2	Short cantilever beam with different loading configurations [3]. (a) Central load. (b) Load at bottom-right corner. (c) Load at top-right corner.	131
4.3	L-shaped structure under different loading configurations [3]. (a) Central load. (b) Load at bottom-right corner.	132
4.4	Optimized material distributions for the cantilever beam under different loading configurations. The layouts obtained using the SIMP method highlight the formation of efficient load-carrying paths, with material concentrated along principal stress directions for each loading case.	134
4.5	Optimized material distributions for the short cantilever beam under different loading configurations. The results demonstrate consistent identification of efficient load paths, even in a reduced design domain.	135
4.6	Optimized material distributions for the L-shaped structure under different loading configurations. The results highlight the ability of the method to capture complex load transfer mechanisms in non-rectangular domains.	137
4.7	Monotonicity-preserving local convex MMA approximations obtained by selectively activating a single reciprocal term in the one-dimensional setting. Depending on the sign of the local derivative $f'(x_0)$, either the upper or the lower moving asymptote governs the curvature and numerical conditioning of the approximation, while first-order consistency with the original nonlinear and nonconvex objective function is preserved at the reference configuration x_0 . (a) Locally increasing objective function, with active upper asymptote ($p > 0, q = 0$). (b) Locally decreasing objective function, with active lower asymptote ($q > 0, p = 0$).	144
4.8	Adaptive update of the moving asymptotes in MMA. The scaling coefficient $s_a^{(k)}$ dynamically modifies the asymptote interval based on the local behavior of the design variable. (a) Oscillatory evolution across successive iterations ($x^{k-2} \rightarrow x^{k-1} \rightarrow x^k$), inducing a contraction of the asymptote interval to enhance numerical stability. (b) Monotonic evolution of the design variable, leading to an expansion of the asymptote interval and promoting faster convergence of the iterative process.	156
4.9	Cantilever beam configuration used in Test Problem 1. The structure consists of five beam elements connected at six nodes. Node 1 is clamped, while a vertical point load is applied at node 6. Each element has a square cross-section of variable height x_a and fixed wall thickness.	158

4.10	Eight-bar truss configuration used in Test Problem 2. The external load is applied at node 5.	159
4.11	Two-bar truss configuration (Test Problem 3). The structure is symmetric, with both bars having identical cross-sectional area x_1 , and base length defined by $2x_2$. An external load \mathbf{F} is applied at node 3.	161
C.1	Architecture of the UQforPy repository. The framework is organized into modular components supporting sampling, global sensitivity analysis, surrogate modeling, Bayesian optimization, Bayesian inversion, and uncertainty quantification. . . .	196
D.1	Predictive responses of different surrogate models for the benchmark functions in Figure 3.2b–d. The solid black line represents the high-fidelity model; fuchsia, purple, and orange dotted lines correspond to LP, LE, and CS models; cyan and blue lines represent GP surrogate models with Matérn and RBF kernels. Red diamonds denote training sample locations. (a) Lévy. (b) Griewank. (c) Forrester.	201
D.2	Mean squared error (MSE) comparison of surrogate models for the benchmark functions in Figure 3.2b–d.	201
E.1	Parameter inference via BI results for the Lévy benchmark. (a) MAP estimates from multiple optimization runs. (b) LS functional profile. (c) NLS functional profile representing the posterior distribution approximation with high-probability region marked by the grey dotted horizontal line representing the 0.95-threshold and highlighted by the coral-coloured portion delimited by the dotted vertical lines.	203
E.2	Parameter inference via BI results for the Griewank benchmark. (a) MAP estimates from multiple initializations. (b) LS functional profile. (c) NLS functional profile revealing non-identifiability due to periodic structure.	203
F.1	Marginal posterior densities obtained from surrogate-assisted MCMC sampling for the Mixed Gaussian–Periodic benchmark. For each chain, the empirical histogram of the samples is shown together with a kernel density estimate. . . .	205
F.2	Superimposed marginal posterior density estimates obtained from all surrogate-assisted MCMC chains. The dominant mode is consistently identified, while secondary modes are sampled with varying frequency.	206
F.3	Grid-based posterior density for the Mixed Gaussian–Periodic benchmark, obtained by evaluating the surrogate-induced likelihood over a dense discretization of the parameter space. The posterior exhibits multiple well-separated modes. .	207

List of Tables

2.1	Summary of the parameters of the EOS M270 machine setup used for the AMBench2018-01 benchmark of the NIST experiment [1].	9
2.2	Mesh strategies adopted for the part-scale PBF-LB/M thermomechanical model of the Inconel 625 cantilever beam. For each mesh strategy, the table reports the number of nodes (nb. nodes), number of elements (nb. elements), number of finite element layers through the cantilever beam height (nb. z -layers), and the approximate time required to complete the high-fidelity thermomechanical numerical simulation.	12
2.3	Temperature-dependent elastic properties of Inconel 625 from <i>ANSYS 2021R2</i> : Young’s modulus and Poisson’s ratio.	13
2.4	Temperature-dependent thermophysical properties of Inconel 625 from <i>ANSYS 2021R2</i> : coefficient of thermal expansion, specific heat, and thermal conductivity.	14
2.5	Machine and process parameters used in the part-scale PBF-LB/M numerical simulations of the Inconel 625 cantilever beam. Values reflect the AMBench2018-01 EOS M270 build, with the dwell time multiplier accounting for sequential printing of the four identical cantilever beams on the build plate.	16
2.6	Uncertainty parameters of the PBF-LB/M printing process considered for the UQ analysis with associated interval.	19
2.7	Properties of the three sparse-grid surrogate models used in this work.	21
2.8	PDF of the parameters resulting from the inverse UQ analysis and used as input for the data-informed forward UQ analysis.	33
2.9	Statistical information for the two uncertain parameters before and after the Bayesian inverse. The reported intervals correspond to approximate 99.7% credible intervals defined as $\mu \pm 3\sigma$	47
3.1	Configuration settings for surrogate modeling across benchmark functions.	68
3.2	MAP estimate and marginal credible intervals for the Mixed Gaussian–Periodic benchmark function. The target value is $\bar{y} = 0.63$, the LS residual at the MAP point is approximately 0.03, and the posterior peak reaches $NLS_{\max} \approx 0.90$	76
4.1	Parameters used in the topology optimization process.	128
4.2	Load values for the 2D cantilever beam under different loading conditions [3].	129

4.3	Load magnitudes for the short cantilever beam under various loading configurations [3].	129
4.4	Load magnitudes for the L-shaped structure under different loading conditions [3].	132
4.5	Comparison of DBF and SBF filtering strategies for the cantilever beam. The compliance values closely match benchmark results, while DBF shows a slightly faster convergence in terms of iteration count.	133
4.6	Comparison of DBF and SBF filtering strategies for the short cantilever beam. The results show excellent agreement with benchmark values and consistent convergence behaviour.	136
4.7	Comparison of DBF and SBF filtering strategies for the L-shaped structure. The results confirm the consistency of both approaches with benchmark data, with minor differences in convergence behaviour.	136
4.8	MMA results for the cantilever beam problem. Each entry reports the pair (weight, infeasibility) at each iteration for different asymptote-spacing parameters t .	165
4.9	Element topology: node connectivity.	165
4.10	Nodal coordinates (in mm).	165
4.11	Optimal cross-sectional areas (in mm) obtained with MMA.	165
4.12	Convergence of total mass (in kg) for Test Problem 2. Results match those reported in [4].	166
4.13	Element topology: node connectivity.	166
4.14	Nodal coordinates (in meters).	166
4.15	Results for Test Problem 3 using MMA. Each row shows x_1 , x_2 , stress σ_1 , and weight w	166

List of Abbreviations

M-AM	Metal Additive Manufacturing
PBF-LB/M	Laser-Based Powder Bed Fusion of Metals
FE	Finite Element
UQ	Uncertainty Quantification
PDF	Probability Density Function
QoI	Quantity of Interest
GSA	Global Sensitivity Analysis
NIST	National Institute of Standards and Technology
MISC	Multi-Index Stochastic Collocation
MCMC	Markov-Chain Monte Carlo
MAP	Maximum A Posteriori
GP	Gaussian Process
AM	Additive Manufacturing
LS	Least Squares
NLS	Negative Least Squares
MSE	Mean Squared Error
LP	Lagrangian Polynomial
LE	Legendre Expansion
CS	Cubic Spline
L-BFGS-B	Limited-memory Broyden Fletcher Goldfarb Shanno with Bounds
FANS	Fourier-Accelerated Nodal Solver
NRMSE	Normalized Root Mean Square Error
NLL	Negative Log-Likelihood
NLP	Negative Log-Posterior
HFM	High-Fidelity Model
BC	Bias Corrected
SIMP	Solid Isotropic Material Penalization
OC	Optimality Criterion
DBF	Density Based Filtering
SBF	Sensitivity Based Filtering

MMA Method of Moving Asymptotes

List of Symbols

Chapter 2: Uncertainty quantification of a part-scale thermomechanical cantilever beam model produced by metal additive manufacturing

Thermal problem

ρ	Constant material density
$c_p(T)$	Temperature-dependent specific heat capacity at constant pressure
T	Temperature field
$k(T)$	Temperature-dependent thermal conductivity
Ω	Physical domain of the part-scale thermomechanical model
t	Time
T_0	Initial uniform temperature
$\partial\Omega$	Boundary of the physical domain
$\partial\Omega_T$	Dirichlet boundary for prescribed temperature
$\partial\Omega_Q$	Neumann boundary for heat flux
$\partial\Omega_H$	Adiabatic boundary (zero heat flux)
\bar{T}	Prescribed environmental temperature
\mathbf{n}	Outward unit normal vector
\bar{q}	Total heat loss through the free surfaces
q_p	Heat loss due to conduction through the powder
q_g	Heat loss due to convection through the environmental gas
h_p	Powder convection coefficient
h_g	Gas convection coefficient
Γ_p	Interface between component and powder
Γ_g	Interface between component and environment gas

Mechanical problem

σ	Cauchy stress tensor
----------	----------------------

\mathbf{D}^{el}	Isotropic elasticity tensor
$\boldsymbol{\varepsilon}^{el}$	Elastic strain tensor
$\boldsymbol{\varepsilon}^{tot}$	Total strain tensor
$\boldsymbol{\varepsilon}^{th}$	Thermal strain tensor
$\boldsymbol{\varepsilon}^{pl}$	Plastic strain tensor
\mathbf{u}	Displacement vector
α	Temperature-dependent thermal expansion coefficient
ΔT	Temperature variation
\mathbf{I}	Second-order identity tensor
$\dot{\boldsymbol{\varepsilon}}^{pl}$	Plastic strain rate
$\dot{\gamma}$	Plastic consistency parameter
Σ	Yield function
σ_{vm}	Equivalent von Mises stress
$\sigma_y(T)$	Temperature-dependent yield stress
\mathbf{s}	Deviatoric stress tensor
$\partial\Omega_U$	Mechanical Dirichlet boundary (fixed support at build plate)
∇	Gradient operator
$\nabla \cdot (\cdot)$	Divergence operator
$\text{tr}(\cdot)$	Trace of a tensor

Numerical Modeling:

nb. nodes	Total number of finite element nodes
nb. elements	Total number of finite elements
nb. z -layers	Number of FE layers through the beam height
T_{melt}	Melting temperature
T_A	Activation temperature assigned to newly activated layers

Uncertainty Quantification:

v_n	Generic uncertain parameter ($n = 1, 2, 3$)
\mathbf{v}	Vector of uncertain parameters
$\rho_n(v_n)$	PDF of parameter v_n
ρ_{prior}	Joint prior PDF of uncertain parameters
$\rho_{\text{prior,red}}$	Reduced prior PDF after GSA
ρ_{post}	Posterior PDF of the uncertain parameters after Bayesian inversion

$\rho_{n,\text{post}}$	Posterior PDF of parameter v_n
$\rho_{\text{prior}}(\mathbf{v})$	Joint prior PDF under independence assumption
Γ_n	Interval of variation of parameter v_n , $[a_n, b_n]$
Γ	Parameter space $\Gamma_1 \times \Gamma_2 \times \Gamma_3$
Γ_{red}	Reduced parameter space $\Gamma_1 \times \Gamma_3$
\mathbf{x}	Generic spatial coordinate in Ω
$\mathbf{x}_{j,\text{GSA}}$	Locations at the centers of the eleven ridges for GSA
$\mathbf{x}_{k,\text{meas}}$	Measurement locations used for inverse UQ
$\mathbf{x}_{j,\text{str}}$	Locations for evaluating residual strains in the forward UQ
$u(\mathbf{x}, \mathbf{v})$	Displacement along the z -direction at point \mathbf{x} depending on \mathbf{v}
$\mathbf{f}(\mathbf{v})$	Vector-valued QoI depending on \mathbf{v}
$f(\mathbf{v})$	Scalar-valued QoI depending on \mathbf{v}
$\varepsilon_{xx}(\mathbf{x}, \mathbf{v})$	Residual strain component at location \mathbf{x}
$\mathcal{S}_{\mathcal{I}}f$	Sparse-grid surrogate model of $f(\mathbf{v})$
$\mathcal{U}_{\mathbf{i}}(\mathbf{v})$	Tensor-product Lagrange interpolant for multi-index \mathbf{i}
\mathcal{I}	Multi-index set used to build sparse-grid surrogate models
\mathcal{I}_{max}	Tensor-grid multi-index set
\mathcal{I}_{sum}	Sparse-grid multi-index set based on summed levels
\mathbf{i}	Multi-index specifying grid refinement per parameter
$c_{\mathbf{i}}$	Combination technique coefficient
N	Number of uncertain parameters
K_n	Number of collocation points along parameter v_n
$m(i_n)$	Level-to-knots function $m(i_n) = 2i_n - 1$
\mathcal{T}_{i_n}	Univariate set of collocation points for parameter v_n
$M_{\mathbf{i}}$	Total number of tensor-grid points, $\prod_{n=1}^N m(i_n)$
$\mathcal{G}_{\mathcal{I}}$	Sparse grid associated with multi-index set \mathcal{I}
w	Level parameter controlling sparse-grid refinement
\tilde{u}_k	Noisy synthetic displacement measurement at $\mathbf{x}_{k,\text{meas}}$
ε_k	Measurement noise realization
σ_{ε_k}	Standard deviation of the measurement noise
$M_k(\mathbf{v})$	Misfit between model prediction and noisy measurement at location k
$\mathcal{L}(\mathbf{v} \tilde{u}_k)$	Likelihood function
C	Normalization constant in Bayes' theorem

$LS(\mathbf{v})$	Least-squares functional $\sum_k M_k^2$
\mathbf{v}_{MAP}	Maximum A Posteriori estimate of \mathbf{v}
J_u	Jacobian matrix of displacements w.r.t. parameters
$H_{\tilde{u}_k}$	Hessian of $u(\mathbf{x}_{k,meas}, \mathbf{v})$ with respect to \mathbf{v}
Σ_{post}	Covariance matrix of Gaussian approximation of ρ_{post}
$\bar{\sigma}$	Standard deviation of noise used in the likelihood
$\bar{\sigma}_{MAP}^2$	Estimate of $\bar{\sigma}^2$ via residuals
E_{PPE}	Pointwise prediction error
E_{MSE}	Root mean square error
M	Number of random samples used in convergence tests
L	Number of strain-evaluation locations
α	Fidelity index of the PBF-LB/M numerical model ($\alpha = 1$ coarse mesh, $\alpha = 2$ fine mesh)
$f_\alpha(\mathbf{v})$	QoI evaluated using fidelity level α .
$\mathcal{T}_m(\boldsymbol{\beta})$	Cartesian grid of collocation points associated with multi-index $\boldsymbol{\beta}$
$\boldsymbol{\beta} = (\beta_1, \dots, \beta_N)$	Multi-index controlling the number of collocation points per dimension
$\mathcal{T}_{n,m}(\beta_n)$	Univariate collocation points for parameter v_n at level $m(\beta_n)$
$\mathbf{v}_m^{(j)}$	Generic collocation point of the Cartesian grid $\mathcal{T}_m(\boldsymbol{\beta})$
$\mathcal{L}_m^{(j)}$	Multivariate Lagrange basis polynomial associated with node $\mathbf{v}_m^{(j)}$
$[\alpha, \boldsymbol{\beta}]$	Extended multi-index identifying fidelity level and stochastic-collocation level
\mathcal{S}_I	MISC surrogate operator combining multiple single-fidelity interpolants
$c_{\alpha,\boldsymbol{\beta}}$	Combination technique coefficient for extended multi-index $[\alpha, \boldsymbol{\beta}]$
\mathcal{J}_u	Jacobian matrix of model predictions w.r.t. uncertain parameters evaluated at \mathbf{v}_{MAP}
S	Number of Monte Carlo samples used in forward UQ analysis
K	Number of displacement measurements used in the Bayesian inversion
J	Number of strain evaluation nodes in the cantilever beam

σ_{meas}	Standard deviation of the measurement noise in the NIST displacement data
$u_{k,exp}$	Experimental displacement measured at location $\mathbf{x}_{k,meas}$
ϵ_k	Experimental noise realization affecting measurement $u_{k,exp}$
\mathcal{L}	Likelihood function for the Bayesian inversion problem
Σ_{post}	Covariance matrix of the Gaussian approximation of the posterior PDF
$\ell_{n,m}^{(j_n)}(\beta_n)$	Univariate Lagrange basis function for parameter v_n
$Q_{prior,\cdot}$	Prior-based quantiles of the predicted residual strain PDFs
$Q_{post,\cdot}$	Posterior-based quantiles of the predicted residual strain PDFs
CoV	Coefficient of Variation (ratio between standard deviation and mean)
MAP	Maximum A Posteriori estimate of the uncertain parameters
S_{Iu}	MISC surrogate approximation of the displacement field
U_Z	Displacement component along the z -axis used in model validation

Chapter 3: Uncertainty-Aware Bayesian Inversion and Adaptive Surrogate Modeling for Inverse Design

An Uncertainty-Aware Framework for Inverse Problems through Adaptive Surrogate Modeling and Bayesian Inversion

\mathbf{x}	Vector of input parameters in the inverse problem
d	Dimension of the parameter space
$f(\mathbf{x})$	High-fidelity forward model mapping parameters to quantity of interest
y	Quantity of interest (model output)
\bar{y}	Observed or target quantity of interest
ϵ	Additive observational noise
σ^2	Variance of the observational noise
$p(\mathbf{x})$	Prior probability density function

$p(\bar{y} \mathbf{x})$	Likelihood function
$p(\mathbf{x} \bar{y})$	Posterior probability density function
\mathbf{x}_{MAP}	Maximum A Posteriori estimate
Σ_{MAP}	Covariance matrix from Laplace approximation at MAP
$f_{\text{GP}}(\mathbf{x})$	Gaussian process surrogate model
$\mu_{\text{GP}}(\mathbf{x})$	Predictive mean of the GP surrogate
$\sigma_{\text{GP}}^2(\mathbf{x})$	Predictive variance of the GP surrogate
$k(\mathbf{x}, \mathbf{x}')$	Covariance kernel of the Gaussian process
D	Training dataset of input–output pairs
n_{init}	Number of initial training samples
n_{val}	Number of validation samples
MSE	Mean squared error used to assess surrogate accuracy
\mathbf{x}_{next}	Next sampling point selected by adaptive strategy
\mathcal{X}	Parameter domain
$\sigma_{\text{GP}}(\mathbf{x})$	Predictive standard deviation used as acquisition criterion
$LS(\mathbf{x})$	Least-squares misfit functional
$NLS(\mathbf{x})$	Negative least-squares functional (unnormalized posterior indicator)
$\ \cdot\ $	Euclidean norm
Γ	Prior covariance matrix
$\mu_{\mathbf{x}}$	Prior mean of the parameters
w	Auxiliary variable in the Lévy function definition

Uncertainty-Aware Bayesian Inversion Using Adaptive Gaussian-Process Surrogate

Modeling for the Inverse Design of Spinodoid Architected Materials

Ω	Reference domain of the microstructure
$\mathbf{x} = (X, Y, Z)^\top$	Spatial coordinate in the reference domain
$\varphi(\mathbf{x})$	Scalar field defining the spinodoid morphology
n_{wave}	Number of wave modes in the random-field construction
\mathbf{k}_k	Wave vector of the k -th harmonic mode
$\hat{\mathbf{k}}$	Normalized wave-vector direction
ϕ_k	Phase shift associated with the k -th wave mode
θ_i	Anisotropy angle defining directional constraints
\mathbf{e}_i	Canonical basis vectors of \mathbb{R}^3

τ	Level-set threshold parameter
Ω_s	Solid domain of the microstructure
Ω_v	Void domain of the microstructure
v_f	Solid volume fraction
$\mathbf{R}(\psi_1, \psi_2, \psi_3)$	Rotation matrix defined by Euler angles
ψ_i	Rotation angles controlling global orientation
Θ	Vector of morphological parameters
S_Θ	Admissible design space
\mathcal{F}	High-fidelity forward operator mapping parameters to responses
\mathbf{y}_{HF}	High-fidelity output vector
\mathbf{C}^{eff}	Effective elasticity tensor
$\boldsymbol{\sigma}^{\text{M}}$	Macroscopic stress tensor
$\boldsymbol{\varepsilon}^{\text{M}}$	Macroscopic strain tensor
E	Young's modulus of the solid phase
ν	Poisson's ratio of the solid phase
K	Bulk modulus
G	Shear modulus
\mathbf{y}	Vector of effective mechanical properties
C_{iii}^{eff}	Diagonal components of the effective stiffness tensor
\mathbf{y}^*	Target response in the inverse problem
\mathcal{D}	Dataset of input–output pairs
N	Number of high-fidelity samples
Θ_i	Input parameter sample
\mathbf{y}_i	Corresponding model output
$\mathcal{X}_{\text{cand}}$	Candidate set for adaptive sampling
N_{cand}	Number of candidate points
$\hat{y}_k(\Theta)$	GP surrogate for the k -th output component
$m_k(\Theta)$	Mean function of the GP prior
$k_k(\Theta, \Theta')$	Covariance kernel of the GP
$\mu_k(\Theta)$	Predictive mean of the GP
$\sigma_k^2(\Theta)$	Predictive variance of the GP
$\boldsymbol{\mu}(\Theta)$	Vector of predictive means
$\boldsymbol{\sigma}(\Theta)$	Vector of predictive standard deviations
$U_{\text{IVAR}}(\Theta)$	Integrated predictive variance criterion
$S(\Theta)$	Adaptive sampling score

β_d	Weighting exponent for distance-based exploration
$S_{AR}(\Theta)$	Posterior-focused refinement score
β_{AR}	Refinement weighting exponent
T	Temperature parameter controlling posterior weighting
$\pi(\Theta)$	Prior distribution
$\pi(\mathbf{y}^* \Theta)$	Likelihood function
$\pi(\Theta \mathbf{y}^*)$	Posterior distribution
$\Sigma_{GP}(\Theta)$	Covariance from GP predictive uncertainty
Σ_{eff}	Effective covariance from residual discrepancy
$\Sigma_{tot}(\Theta)$	Total covariance matrix
$\mathcal{L}(\Theta)$	Negative log-likelihood
Θ_{MAP}	Maximum A Posteriori estimate
$\mathcal{P}(\Theta)$	Additional regularization or penalty term
$\mu_\delta(\Theta)$	Bias correction term
$\mu_{BC}(\Theta)$	Bias-corrected predictive mean

Chapter 4: Structural optimization strategies with optimality criteria and method of moving asymptotes on two- and three-dimensional benchmarks

Solid Isotropic Material Penalization Method

Ω	Design domain of the topology optimization problem
$\partial\Omega$	Boundary of the design domain
Γ_D	Dirichlet boundary where displacements are prescribed
Γ_N	Neumann boundary where tractions are prescribed
\mathbf{x}	Generic point in the design domain
\mathbf{y}	Neighbouring point used in filtering/regularization operators
ϕ	SIMP design variable (material density field), with $0 \leq \phi \leq 1$
$\phi = 0$	Void (absence of material)
$\phi = 1$	Solid material
p	Penalization exponent in SIMP interpolation
$\mathbb{C}(\phi)$	Penalized (effective) stiffness tensor depending on ϕ
\mathbb{C}_0	Elastic stiffness tensor of the solid material
\mathbf{u}	Displacement field
σ	Cauchy stress tensor

$\boldsymbol{\varepsilon}$	Small-strain tensor
$\nabla \mathbf{u}$	Displacement gradient
$\nabla^s \mathbf{u}$	Symmetric gradient of \mathbf{u} , $\nabla^s \mathbf{u} = (\nabla \mathbf{u} + (\nabla \mathbf{u})^T)/2$
\mathbf{b}	Body force per unit volume
\mathbf{t}	Prescribed traction vector on Γ_N
\mathbf{n}	Outward unit normal vector on $\partial\Omega$
$\text{div}(\cdot)$	Divergence operator
:	Double contraction between tensors
$\bar{\mathbf{u}}$	Prescribed displacement on Γ_D
$\delta \mathbf{p}$	Admissible virtual displacement (test function) satisfying Dirichlet constraints
$S(\phi, \mathbf{u}, \delta \mathbf{p})$	Weak equilibrium residual functional (variational form of equilibrium)
$d\Omega$	Area/volume measure in Ω (depending on dimension)
$d\Gamma$	Boundary measure on Γ_N
$C(\phi, \mathbf{u})$	Compliance functional (work of external loads)
$\mathbf{u}(\phi)$	Equilibrium displacement field induced by the density distribution ϕ
$\mathcal{V}(\phi)$	Volume (material) constraint functional.
V	Reference volume of the design domain (e.g., $V = \int_{\Omega} 1 \, d\Omega$)
f	Prescribed volume fraction, $f \in (0, 1]$
$\mathcal{L}(\phi, \mathbf{u}, \lambda, \mathbf{p})$	Lagrangian functional combining objective, constraints, and equilibrium
λ	Lagrange multiplier associated with the volume constraint
\mathbf{p}	Adjoint variable associated with the equilibrium constraint
$\mathcal{S}(\phi, \mathbf{u}, \mathbf{p})$	Lagrangian term enforcing equilibrium through the adjoint field
$\mathcal{L}_{\text{eq}}(\phi, \lambda)$	Equilibrated Lagrangian obtained after enforcing equilibrium (self-adjoint case)
$C_{\text{eq}}(\phi)$	Compliance evaluated at equilibrium, i.e., $C(\phi, \mathbf{u}(\phi))$
$\frac{\partial C(\phi)}{\partial \phi}$	Sensitivity of the penalized stiffness tensor with respect to ϕ

B_{eq}	OC control variable measuring deviation from pointwise optimality ($B_{\text{eq}} = 1$ at optimality)
$\frac{\partial C_{\text{eq}}}{\partial \phi}$	Compliance sensitivity with respect to the density field
$\frac{\partial \mathcal{V}}{\partial \phi}$	Volume sensitivity with respect to the density field
\mathcal{K}	Generic filtering operator used to regularize densities or sensitivities
$\kappa(\cdot)$	Filter kernel function (e.g., RBF compactly supported kernel)
$r(\mathbf{x}, \mathbf{y})$	Normalized distance between points \mathbf{x} and \mathbf{y}
r_{min}	Filter radius (minimum length scale controlling the neighbourhood influence)
$r_0(\mathbf{x}, \mathbf{y})$	Compact-support factor, $r_0 = \max(0, 1 - r)$
$\tilde{\phi}(\mathbf{x})$	Filtered (regularized) density field (notation may be used for $\mathcal{K}(\phi)$)
$\mathcal{K}\left(\frac{\partial C_{\text{eq}}}{\partial \phi}\right)$	Sensitivity filtering operator applied to compliance sensitivities
Ω_e	Domain of finite element e
e	Index of a finite element in the mesh
n_e	Total number of finite elements in the mesh
ϕ_e	Element-wise density (design variable assumed constant over element e)
ϕ_{min}	Minimum density used to avoid singular stiffness in void regions
\mathbf{N}_u	Matrix/vector of shape functions for displacement interpolation
\mathbf{u}_e	Element displacement field (approximated within element e)
\mathbf{d}_e	Vector of nodal displacement degrees of freedom of element e
\mathbf{d}	Global nodal displacement vector (assembled over the whole mesh)
\mathbf{B}_e	Strain–displacement matrix of element e (derivatives of shape functions)
$\boldsymbol{\varepsilon}_e$	Element strain vector/tensor computed as $\boldsymbol{\varepsilon}_e = \mathbf{B}_e \mathbf{d}_e$.
$\boldsymbol{\sigma}_e$	Element stress vector/tensor computed via constitutive law
$\mathbf{K}_e(\phi_e)$	Element stiffness matrix depending on ϕ_e

$\mathbf{K}(\phi)$	Global stiffness matrix assembled from $\mathbf{K}_e(\phi_e)$
\mathbf{f}	Global force vector (assembled external and body forces)
$\mathbf{K}^T = \mathbf{K}$	Symmetry of the stiffness matrix (self-adjoint setting)
$\mathbb{C}(\phi)$	Penalized (effective) stiffness tensor in SIMP
$\mathbb{C}(\phi_e)$	Element stiffness tensor evaluated at element density ϕ_e
\mathbb{C}_0	Stiffness tensor of the solid material
p	Penalization exponent in the discrete SIMP law
$\mathbf{Kd} = \mathbf{f}$	Assembled linear system enforcing discrete equilibrium
$c(\phi, \mathbf{d})$	Discrete compliance objective, $c = \mathbf{f}^T \mathbf{d}$
$c_{\text{eq}}(\phi)$	Compliance evaluated at equilibrium, i.e., $c(\phi, \mathbf{d}(\phi))$
$v(\phi)$	Discrete volume constraint functional
V_e	Volume (or area in 2D) of element e
V_0	Total initial design volume (or area in 2D)
f	Prescribed volume fraction, $f \in (0, 1]$
$\mathcal{L}_{\text{eq}}(\phi, \lambda)$	Equilibrium Lagrangian for the discrete problem, $\mathcal{L}_{\text{eq}} = c_{\text{eq}}(\phi) + \lambda v(\phi)$
λ	Lagrange multiplier associated with the volume constraint
$\frac{d\mathbf{d}}{d\phi}$	Sensitivity of displacements with respect to density variables
$\frac{d\mathbf{K}}{d\phi}$	Sensitivity of global stiffness matrix with respect to density variables
$\frac{d\mathbf{f}}{d\phi}$	Sensitivity of load vector with respect to density variables (often zero)
$\frac{dc_{\text{eq}}}{d\phi}$	Total derivative of equilibrated compliance with respect to density variables
k	Iteration counter in the OC algorithm
ϕ^k	Density field at iteration k
λ^{k+1}	Updated Lagrange multiplier at iteration $k + 1$ (typically via bisection)
η	Damping exponent in OC update (commonly $\eta = 0.5$)

m	Move limit controlling the maximum change in ϕ per iteration
$B_{\text{eq}}(\phi^k, \lambda^{k+1})$	OC control variable (ratio of compliance and volume sensitivities)
$\phi_{\text{trial}}^{k+1}$	Trial density computed from OC update before move-limit and projection
ϕ_{aux}	Auxiliary density after enforcing move limits
ϕ_{project}	Projected density after enforcing bounds $[0, 1]$
$\mathcal{N}(e)$	Neighbourhood set of elements surrounding element e
\mathbf{x}_e	Centroid (representative point) of element e
\mathbf{x}_j	Centroid (representative point) of neighbouring element j
$r(\mathbf{x}_e, \mathbf{x}_j)$	Normalized distance between element centroids, $r = \ \mathbf{x}_e - \mathbf{x}_j\ /r_{\text{min}}$
r_{min}	Filter radius defining the influence distance in the discrete filter
$r_0(\mathbf{x}_e, \mathbf{x}_j)$	Compact-support factor, $r_0 = \max(0, 1 - r)$
$\kappa(r)$	Compactly supported RBF kernel used for filtering
ϕ_e^{k+1}	Filtered density of element e at iteration $k + 1$ (weighted average over $\mathcal{N}(e)$)
$\frac{\partial c_{\text{eq}}}{\partial \phi_e}$	Element-wise compliance sensitivity (possibly filtered)

Method of Moving Asymptotes

$\mathbf{x} = [x_1, \dots, x_A]^T$	Vector of design variables
x_i	Generic design variable
A	Number (dimension) of design variables
$\mathbb{R}_{\text{adm}}^A$	Admissible design set (box-constrained domain) for \mathbf{x}
$\mathbf{x}^{\text{min}}, \mathbf{x}^{\text{max}}$	Lower and upper bounds defining $\mathbb{R}_{\text{adm}}^A$
$f(\mathbf{x})$	Objective function of the constrained optimization problem
$\mathbf{h}(\mathbf{x})$	Vector of equality constraints
$\mathbf{g}(\mathbf{x})$	Vector of inequality constraints
B	Number (dimension) of equality constraints
C	Number (dimension) of inequality constraints
$\nabla f(\mathbf{x})$	Gradient of f with respect to \mathbf{x}

$\nabla^2 f(\mathbf{x})$	Hessian matrix of f with respect to \mathbf{x}
$\nabla \mathbf{g}(\mathbf{x})$	Jacobian matrix of \mathbf{g} with respect to \mathbf{x}
$\nabla \mathbf{h}(\mathbf{x})$	Jacobian matrix of \mathbf{h} with respect to \mathbf{x}
\mathbf{x}^*	Generic stationary point of the objective (or of a Lagrangian formulation)
$\bar{\mathbf{x}}$	Reference (current) design configuration used to build local approximations
$\tilde{f}(\mathbf{x}; \bar{\mathbf{x}})$	Local approximation of the objective function about $\bar{\mathbf{x}}$
$\tilde{\mathbf{g}}(\mathbf{x}; \bar{\mathbf{x}})$	Local approximation of inequality constraints about $\bar{\mathbf{x}}$
$\tilde{\mathbf{h}}(\mathbf{x}; \bar{\mathbf{x}})$	Local approximation of equality constraints about $\bar{\mathbf{x}}$
$\rho_f(\bar{\mathbf{x}})$	Constant term in the separable approximation of the objective
$\boldsymbol{\rho}_g(\bar{\mathbf{x}})$	Constant vector in the separable approximation of \mathbf{g}
$\boldsymbol{\rho}_h(\bar{\mathbf{x}})$	Constant vector in the separable approximation of \mathbf{h}
$\varphi_i(x_i; \bar{\mathbf{x}})$	Univariate component of the separable approximation of \tilde{f}
$\psi_i(x_i; \bar{\mathbf{x}})$	Univariate vector component of the separable approximation of $\tilde{\mathbf{g}}$
$\chi_i(x_i; \bar{\mathbf{x}})$	Univariate vector component of the separable approximation of $\tilde{\mathbf{h}}$
$m_{f,i}(\bar{\mathbf{x}})$	Curvature lower bound enforcing strict convexity of φ_i on $[x_i^{\min}, x_i^{\max}]$
$\mathbf{m}_{g,i}(\bar{\mathbf{x}})$	Vector curvature bounds enforcing convexity of ψ_i (inequality constraints)
$\mathbf{m}_{h,i}(\bar{\mathbf{x}})$	Vector curvature bounds enforcing convexity of χ_i (equality constraints)
$\text{diag}(\cdot)$	Diagonal matrix built from the input vector entries
$\mathbf{0}$	Zero vector (dimension implied by context)
$\mathbf{0}$ (matrix)	Zero matrix (dimension implied by context)
$\mathcal{P}(\bar{\mathbf{x}})$	Local convex separable subproblem associated with the reference point $\bar{\mathbf{x}}$
L	Lower moving asymptote defining the left boundary of the approximation interval
U	Upper moving asymptote defining the right boundary of the approximation interval
(L, U)	Moving-asymptotes interval defining the analytical domain of the convex approximation

- p Positive curvature parameter associated with the upper reciprocal term $p/(U - x)$
- q Positive curvature parameter associated with the lower reciprocal term $q/(x - L)$
- r Constant shift ensuring value consistency between $\tilde{f}(x)$ and $f(x)$ at x_0
- $\tilde{f}'(x)$ First derivative of the convex MMA approximation
- $\tilde{f}''(x)$ Second derivative (curvature) of the convex MMA approximation
- κ Prescribed positive curvature used for regularization when $f'(x_0) = 0$
- α Curvature-splitting parameter distributing κ between reciprocal terms
- σ Baseline half-width of the symmetric moving-asymptotes interval
- τ_{\min} Minimum admissible distance between the reference point x_0 and each asymptote
- $\delta(x_0)$ Local variation indicator controlling adaptive spacing of the moving asymptotes
- s Scaling coefficient governing the adaptation rate of the moving asymptotes
- $\mathcal{O}((x - x_0)^2)$ Second-order local approximation error of the MMA surrogate near x_0
- $g(x)$ Nonlinear and possibly nonconvex inequality constraint function in the one-dimensional structural optimization problem
- $\tilde{g}(x)$ Local convex reciprocal approximation of the constraint function $g(x)$ within the moving-asymptote interval (L, U)
- r_g Constant (offset) parameter of the convex approximation of the constraint function
- p_g Positive curvature parameter associated with the reciprocal term $1/(U - x)$ in the convex approximation of $g(x)$
- q_g Positive curvature parameter associated with the reciprocal term $1/(x - L)$ in the convex approximation of $g(x)$

- $\mathcal{L}(x, \lambda)$ Lagrangian functional associated with the one-dimensional convex subproblem, combining objective and constraint approximations
- λ Lagrange multiplier associated with the inequality constraint in the MMA convex subproblem
- (x^*, λ^*) Optimal primal–dual solution pair satisfying the KKT conditions of the convex subproblem
- $\partial \mathcal{L} / \partial x$ Derivative of the Lagrangian functional with respect to the primal design variable x
- A Number of design variables in the multidimensional optimization problem
- x_a a -th design variable, $a = 1, \dots, A$
- \mathbf{x}_0 Reference design configuration at which the local convex approximation is constructed
- x_{0a} a -th component of the reference configuration \mathbf{x}_0
- L_a Lower moving asymptote associated with design variable x_a
- U_a Upper moving asymptote associated with design variable x_a
- p_a Curvature parameter of the convex MMA approximation associated with $(U_a - x_a)^{-1}$
- q_a Curvature parameter of the convex MMA approximation associated with $(x_a - L_a)^{-1}$
- $f_{x_a}(\mathbf{x}_0)$ Partial derivative of the objective function with respect to x_a , evaluated at \mathbf{x}_0
- κ_a Prescribed positive curvature value for the MMA approximation in direction x_a when $f_{x_a}(\mathbf{x}_0) = 0$
- α_a Weighting parameter controlling curvature distribution between the reciprocal terms in direction x_a
- $\mathbf{p}_{g,a}$ Vector of curvature parameters for constraint approximations associated with $(U_a - x_a)^{-1}$
- $\mathbf{q}_{g,a}$ Vector of curvature parameters for constraint approximations associated with $(x_a - L_a)^{-1}$
- $p_{g_{c,a}}$ Curvature parameter of the c -th constraint approximation in direction x_a for $(U_a - x_a)^{-1}$
- $q_{g_{c,a}}$ Curvature parameter of the c -th constraint approximation in direction x_a for $(x_a - L_a)^{-1}$

- $g_{c,x_a}(\mathbf{x}_0)$ Partial derivative of constraint g_c with respect to x_a , evaluated at \mathbf{x}_0
- $\tau_{\min,a}$ Minimum admissible separation between x_{0a} and the moving asymptotes L_a and U_a

Chapter 1

Introduction

The present chapter provides an overview of the doctoral thesis by outlining the motivations that drive the research programme (Section 1.1), defining the objectives that guide the investigation (Section 1.2), and describing the structure and organization of the thesis (Section 1.3).

1.1 Motivations

Engineering decision-making in computational mechanics requires the systematic treatment of uncertainty, high computational cost, and complex multiphysics interactions. Substantial research efforts have focused on the development of uncertainty quantification methodologies for large-scale engineering systems, including stochastic finite element methods, Bayesian inference, and surrogate-based approaches [5–8]. These methodologies have been increasingly applied in application-driven contexts such as metal additive manufacturing, architected materials, and structural optimization, where process–structure–property relationships govern system behavior and design performance [9–11]. In these settings, predictive models must remain reliable under variability in material properties, process parameters, and operating conditions, while maintaining computational tractability.

In metal additive manufacturing [12, 13], multiscale thermomechanical phenomena evolve under multiple sources of uncertainty affecting both process behavior and predictive modeling. Melt-pool dynamics, heat transfer, phase transformations, and residual strain accumulation exhibit strong sensitivity to process parameters, material properties, and boundary conditions [9, 14–16]. These mechanisms govern microstructural evolution, anisotropy, residual stresses, and geometric distortion, and therefore directly affect the structural performance of manufactured components. However, experimental observations are often limited, indirect, or affected by measurement uncertainty, which complicates model calibration and validation [17, 18]. Recent developments in high-fidelity numerical modeling and data-driven approaches highlight the need for methodologies that enable efficient parameter identification and uncertainty-aware predictions [19, 20]. Despite these advances, high computational cost and the difficulty of incorporating uncertainty in a consistent manner remain significant challenges.

Architected materials such as spinodoids introduce additional challenges associated with high-dimensional design spaces and complex structure–property relationships [21–24]. The mapping between geometric parameters and effective mechanical properties is typically nonlinear, nonconvex, and highly sensitive to perturbations, which limits the applicability of exhaustive

numerical exploration strategies. Computational homogenization provides a rigorous framework for evaluating effective properties [25, 26], but its computational cost motivates the use of surrogate modeling techniques [27, 28]. Gaussian-process surrogate models enable probabilistic approximations that provide both predictive mean and uncertainty, supporting adaptive sampling strategies and inverse-design formulations [29, 30]. However, surrogate-based approaches introduce approximation errors and require a consistent treatment of predictive uncertainty and model discrepancy, particularly in inverse design settings.

Structural topology optimization [11, 31, 32] represents a further domain in which modeling complexity and numerical challenges play a central role. Density-based approaches such as the Solid Isotropic Material Penalization (SIMP) method and gradient-based optimization strategies such as the Method of Moving Asymptotes (MMA) enable the solution of large-scale design problems governed by partial differential equations [4, 33]. The governing equations depend on a material distribution that evolves during the optimization process and must remain consistent with nonlinear material behavior, numerical stability requirements, and manufacturability constraints. The resulting optimization problems are typically nonconvex and may exhibit multiple local minima, mesh-dependence, and sensitivity to regularization and algorithmic parameters [34, 35]. Despite their effectiveness, these methods require careful parameter tuning and their extension to uncertainty-aware settings remains limited.

Across these domains, common challenges emerge. High-fidelity numerical models require substantial computational resources, limiting systematic exploration of parameter and design spaces. Uncertainty arises from variability in inputs, incomplete physical modeling, and model-form discrepancies; without explicit treatment, these factors reduce predictive reliability and may lead to biased or non-robust design decisions [6, 30]. Inverse problems associated with parameter identification are often ill-posed and sensitive to limited or noisy data, motivating probabilistic formulations that explicitly characterize uncertainty and support regularized inference [36–38].

1.2 Objectives

The considerations outlined above motivate three complementary but distinct research directions addressing modeling, computational, and uncertainty-related challenges in metal additive manufacturing, architected materials, and structural optimization.

The first research direction develops an uncertainty-quantification framework for part-scale thermomechanical modeling in metal additive manufacturing. The framework combines variance-based global sensitivity analysis, Bayesian inverse calibration, and data-informed forward uncertainty propagation. The objective consists in quantifying the impact of variability in process parameters on quantities of interest, identifying dominant sources of uncertainty, and enabling model calibration under limited data. High-dimensional responses are approximated using single- and multi-fidelity sparse-grid surrogate models, enabling efficient uncertainty analysis while controlling computational cost.

The second research direction investigates inverse problems and inverse design through Bayesian approaches combined with surrogate modeling. Gaussian-process surrogate models

are constructed and adaptively refined based on predictive uncertainty, enabling efficient approximation of computationally expensive forward models. The objective consists in enabling computationally efficient inversion while maintaining a consistent probabilistic treatment of uncertainty. The inverse problem is formulated within a Bayesian framework and addressed through approximate posterior representations based on local Laplace approximations around the maximum a posteriori estimate. Particular attention is devoted to the consistent treatment of surrogate-induced uncertainty and model discrepancy within the inversion process.

The third research direction examines structural topology and nonlinear optimization methods, with particular emphasis on the Optimality Criteria method and the Method of Moving Asymptotes. The objective consists in providing a rigorous and unified analysis of these established methodologies, clarifying their theoretical structure, convergence properties, and numerical behavior on benchmark problems, and establishing a methodological basis for future developments in uncertainty-aware optimization.

Across the three research directions, the common objective is to develop methodologies that reduce computational cost, explicitly represent uncertainty when relevant, and support reliable interpretation of model predictions in complex engineering systems.

Although the three directions are developed independently, they share common methodological components, including surrogate modeling, approximation strategies, and iterative solution procedures. Their integration within a unified uncertainty-aware design framework, enabling the consistent propagation of uncertainty from forward modeling to inverse analysis and design optimization, represents a natural and promising direction for future research.

1.3 Structure and Organization of the Thesis

The thesis is organized as follows. Chapter 2 develops the uncertainty-quantification framework for thermomechanical modeling in metal additive manufacturing. Chapter 3 presents surrogate-based Bayesian methodologies for inverse problems and inverse design, including adaptive surrogate modeling and applications to architected materials. Chapter 4 provides a theoretical and numerical analysis of structural topology optimization and nonlinear optimization methods, focusing on the SIMP approach with Optimality Criteria and on the Method of Moving Asymptotes. Chapter 5 concludes the thesis by summarizing the main contributions and outlining future research directions, with particular emphasis on potential synergies between the different methodologies.

Chapter 2

Uncertainty quantification of a part-scale thermomechanical cantilever beam model produced by metal additive manufacturing

2.1 Introduction

Metal Additive Manufacturing (M-AM) [12, 39] encompasses a family of advanced manufacturing technologies capable of producing components with complex geometries, tailored internal architectures, and high mechanical performance through a layer-by-layer fabrication strategy. Among the existing M-AM technologies for metal materials [13], Laser-Based Powder Bed Fusion of Metals (PBF-LB/M) represents one of the most widely adopted processes in industrial applications [14, 40]. In PBF-LB/M, each fabrication cycle begins with the deposition of a thin layer of metal powder across a build plate inside a sealed process chamber filled with an inert gas, typically argon or nitrogen, to limit oxidation and ensure stable thermal conditions. A high-energy-density laser beam selectively melts the powder following a predefined scan strategy dictated by the part geometry. After solidification of the molten track, a new powder layer is deposited and the process is repeated until completion of the component.

Despite significant technological progress, the process–structure–property relationships governing PBF-LB/M remain only partially characterized due to the strongly coupled multi-physics and multi-scale nature of the process [41]. Rapid melting and solidification, heterogeneous thermal gradients, melt-pool instabilities, microstructural evolution, and residual stress accumulation contribute to a complex manufacturing environment in which reliable parameter identification remains challenging. As a consequence, calibration of process and material parameters often relies on costly experimental campaigns. In this context, computer-aided numerical simulations have become an essential tool for analyzing, optimizing, and validating M-AM processes [42, 43]. Numerical models for PBF-LB/M span several characteristic length scales and are typically categorized into microscale, mesoscale, and macroscale models [44].

Microscale (powder-scale) models explicitly resolve powder particles, laser–material interaction, melt-pool solidification, and grain evolution. Such models provide detailed insight into microstructural changes and are commonly formulated using Phase-Field methods [43, 45–47] or Cellular Automata [48]. Mesoscale models focus on melt-pool dynamics, thermal–fluid flow, recoil pressure, evaporation phenomena, and scanning strategy effects. Representative approaches include the Lattice-Boltzmann method [49, 50] and the Discrete Element method [51, 52], en-

abling the analysis of melt-pool geometry, defect formation, and track stability. Macroscale or part-scale models, typically based on the Finite Element method [15, 16, 46, 53–57], describe the global thermomechanical response of the component, including transient temperature fields, residual stresses, and distortions.

PBF-LB/M processes inherently involve multiple sources of uncertainty, arising from powder characteristics (particle size distribution, morphology, packing density), material properties (thermal conductivity, heat capacity, absorptivity), boundary conditions, and operational parameters such as laser power, scanning speed, hatch spacing, and inter-layer dwell times [14]. Uncertainty Quantification (UQ) methodologies [58, 59] provide a rigorous framework for quantifying, propagating, and mitigating the influence of these uncertainties on both process behavior and resulting product performance. Forward UQ propagates uncertainties from input parameters — modeled as random variables equipped with probability density functions (PDFs) — to model outputs, commonly referred to as Quantities of Interest (QoIs). Inverse UQ reduces parameter uncertainty by identifying data-informed posterior PDFs consistent with available measurements, typically within a Bayesian inference framework [60–62]. In contrast to deterministic calibration, Bayesian approaches provide a probabilistic characterization of parameter uncertainty.

Although UQ methodologies have been successfully applied in experimental studies and reduced-order models of M-AM processes [59, 63, 64], their application to high-fidelity numerical models of PBF-LB/M remains challenging due to the computational cost associated with repeated model evaluations. Several studies have addressed forward UQ through uncertainty propagation analyses [65], process monitoring and anomaly detection [66], stochastic modeling of microstructural variability [67], and multi-level representations of material uncertainty [68]. However, inverse UQ at the part scale remains significantly more demanding and comparatively less explored.

To alleviate the computational burden associated with UQ, surrogate modeling has emerged as a key enabling approach. Surrogate models approximate the input–output mapping of high-fidelity models using a limited set of simulations, thereby enabling efficient probabilistic analyses. In the context of M-AM, Gaussian process surrogate models have been widely adopted. Representative contributions include Bayesian quantification of model and data uncertainties in melt-pool geometry prediction [69, 70], uncertainty propagation from process parameters to microstructure and mechanical properties [71], sequential Bayesian calibration strategies [72, 73], surrogate-assisted multiscale modeling of solidification microstructures [68, 74], and probabilistic characterization of micro-segregation [75]. These contributions highlight the role of surrogate modeling in enabling tractable uncertainty analysis for complex PBF-LB/M processes.

The present chapter builds upon two methodological contributions previously developed by the author [76, 77], which together establish a computationally efficient UQ framework for part-scale thermomechanical modeling of PBF-LB/M. The first contribution introduces a workflow that integrates variance-based Global Sensitivity Analysis (GSA), Bayesian inverse UQ, and data-informed forward UQ for the AMBench2018-01 Inconel 625 cantilever benchmark provided by the National Institute of Standards and Technology (NIST) [1, 2]. The workflow identifies the most influential thermomechanical parameters, constructs data-informed PDFs for activation temperature and convection coefficients based on displacement measurements,

and quantifies the resulting uncertainty in residual strain predictions. Single-fidelity sparse-grid surrogate models based on stochastic collocation within the Smolyak framework [78–80] enable accurate approximation of high-dimensional responses with a limited number of simulations.

The second contribution extends the framework by introducing a multi-fidelity sparse-grid surrogate modeling strategy based on the Multi-Index Stochastic Collocation (MISC) method [81, 82]. Low-resolution simulations capture the global parameter dependence, while selected high-resolution simulations correct modeling discrepancies. The resulting multi-fidelity surrogate models enable accurate posterior inference and forward uncertainty propagation at reduced computational cost. The probabilistic predictions show reduced uncertainty compared to prior-based analyses and exhibit agreement with experimental displacement and residual strain data provided by NIST [1, 2].

The developments presented in this chapter provide a structured methodology for uncertainty quantification in thermomechanical modeling of PBF-LB/M processes. The chapter is organized as follows. Section 2.2 presents the governing equations of the thermal and mechanical problems. Section 2.3 describes the AMBench2018-01 experimental setup. Section 2.4 details the numerical modeling strategy. Section 2.5 presents the UQ framework and corresponding results. Finally, Section 2.6 summarizes the main findings and outlines possible directions for further investigation.

2.2 Governing equations

The section introduces the governing thermal and mechanical equations underlying the part-scale thermomechanical FE model adopted for the numerical simulation of the PBF-LB/M process [15, 16, 42, 53–56, 83, 84].

2.2.1 Thermal problem

Assuming that the material follows Fourier’s law of heat conduction, the heat transfer equation takes the form [85]:

$$\rho c_p(T) \frac{\partial T}{\partial t} - \nabla \cdot (k(T) \nabla T) = 0 \quad \text{in } \Omega, \quad (2.1)$$

where T denotes the temperature field, ρ is the (assumed constant) material density, $c_p(T)$ is the temperature-dependent specific heat capacity at constant pressure, $k(T)$ is the temperature-dependent thermal conductivity, and Ω is the physical domain of the part-scale thermomechanical model.

The initial condition at time $t = 0$ prescribes a uniform temperature distribution as:

$$T = T_0 \quad \text{in } \Omega, \quad (2.2)$$

while the boundary $\partial\Omega = \partial\Omega_Q \cup \partial\Omega_T \cup \partial\Omega_H$ is subject to Dirichlet, Neumann, and adiabatic

conditions:

$$T = \bar{T} \quad \text{on } \partial\Omega_T \subset \partial\Omega, \quad (2.3)$$

$$-k(T) \nabla T \cdot \mathbf{n} = \bar{q} \quad \text{on } \partial\Omega_Q \subset \partial\Omega, \quad (2.4)$$

where \bar{T} denotes the prescribed environmental temperature on the lateral surface of the build plate, whereas \bar{q} represents the heat loss through the free surfaces with outward unit normal \mathbf{n} . Adiabatic conditions are imposed on the remaining portion of the boundary, $\partial\Omega_H$.

In part-scale numerical simulations of the PBF-LB/M process, the heat loss through $\partial\Omega_Q$ originates from two physical mechanisms: conduction through the powder, denoted by q_p , and convection through the environmental gas, denoted by q_g . The total heat loss therefore satisfies:

$$\bar{q} = q_p + q_g. \quad (2.5)$$

Since the powder is not explicitly modeled, both heat-loss mechanisms are represented through convection-like boundary conditions. Two heat transfer coefficients are introduced: a powder convection coefficient h_p and a gas convection coefficient h_g . This modeling choice follows common practice in part-scale numerical simulations, where it is generally difficult to distinguish between powder- and gas-driven heat transfer phenomena [15]. According to Newton's cooling law, the two contributions take the form:

$$q_p = h_p (T - \bar{T}) \quad \text{on } \Gamma_p \subset \partial\Omega_Q, \quad (2.6)$$

$$q_g = h_g (T - \bar{T}) \quad \text{on } \Gamma_g \subset \partial\Omega_Q, \quad (2.7)$$

where Γ_p and Γ_g denote the interfaces between the component and the powder and environmental gas, respectively. The interfaces between the component and the powder and the environmental gas satisfy:

$$\partial\Omega_Q = \Gamma_p \cup \Gamma_g, \quad \Gamma_p \cap \Gamma_g = \emptyset.$$

2.2.2 Mechanical problem

The mechanical response of the component is governed by the equilibrium equation:

$$\nabla \cdot \boldsymbol{\sigma} = \mathbf{0} \quad \text{in } \Omega, \quad (2.8)$$

where $\boldsymbol{\sigma}$ denotes the Cauchy stress tensor. A thermo-elasto-plastic constitutive model is adopted, namely:

$$\boldsymbol{\sigma} = \mathbf{D}^{el} \boldsymbol{\varepsilon}^{el}, \quad (2.9)$$

where \mathbf{D}^{el} is the isotropic elasticity tensor, expressed in terms of the Young's modulus and Poisson's ratio, and $\boldsymbol{\varepsilon}^{el}$ is the elastic strain tensor.

The total strain $\boldsymbol{\varepsilon}^{tot}$ is decomposed into elastic, thermal, and plastic contributions:

$$\boldsymbol{\varepsilon}^{tot} = \boldsymbol{\varepsilon}^{el} + \boldsymbol{\varepsilon}^{th} + \boldsymbol{\varepsilon}^{pl} = \frac{1}{2} [\nabla \mathbf{u} + (\nabla \mathbf{u})^T], \quad (2.10)$$

with \mathbf{u} the displacement vector.

The thermal strain acts as an external thermomechanical load and is defined as:

$$\boldsymbol{\varepsilon}^{th} = \alpha \Delta T \mathbf{I}, \quad (2.11)$$

where $\alpha = \alpha(T)$ denotes the temperature-dependent thermal expansion coefficient, $\Delta T = T - T_0$ is the temperature variation, and \mathbf{I} is the second-order identity tensor.

Plastic deformation follows a rate-dependent Prandtl-Reuss flow rule [15, 83]:

$$\dot{\boldsymbol{\varepsilon}}^{pl} = \dot{\gamma} \frac{\partial \Sigma}{\partial \boldsymbol{\sigma}}, \quad (2.12)$$

where $\dot{\gamma}$ is the plastic consistency parameter, and Σ is the yield function:

$$\Sigma = \sigma_{vm} - \sigma_y \leq 0.$$

The equivalent von Mises stress is given by:

$$\sigma_{vm} = \sqrt{\frac{3}{2} \mathbf{s} : \mathbf{s}}, \quad \mathbf{s} = \boldsymbol{\sigma} - \frac{1}{3} \text{tr}(\boldsymbol{\sigma}) \mathbf{I},$$

and the yield stress $\sigma_y = \sigma_y(T)$ depends on temperature.

The mechanical problem is supplemented with the Dirichlet boundary condition:

$$\mathbf{u} = \mathbf{0} \quad \text{on } \partial\Omega_U \subset \partial\Omega, \quad (2.13)$$

where $\partial\Omega_U$ corresponds to the bottom surface of the build plate, representing a fixed support. Homogeneous Neumann boundary conditions are prescribed on the remaining portion of $\partial\Omega$.

2.3 Experimental setup

The present section describes the AMBench2018-01 benchmark experiment performed by NIST [1, 2], which serves as the reference for the present study. The geometry of the experiment corresponds to a bridge structure, shown in Figure 2.1. All information related to the design and setup of the experiment is available on the AMBench2018 website [1].

As shown in Figure 2.1, four identical metal parts (cantilever beams) were constructed on the build plate for each construction process. The study considers only one cantilever beam since the four structures are spaced sufficiently far apart to reasonably neglect thermal and mechanical interaction. However, to correctly capture the cooling time between layers, the entire build layout must be taken into account. For this reason, the model adopted in the present work (see Section 2.4) reproduces a single cantilever beam built on a build plate measuring 85 mm × 12 mm × 20 mm (length × height × width).

According to [1], the experiment was carried out using both 15-5 stainless steel and Inconel 625 superalloy. The present work focuses exclusively on the latter alloy, whose material properties




Figure 2.1: AMBench2018-01 benchmark for the experiment performed by NIST [1, 2].

Table 2.1: Summary of the parameters of the EOS M270 machine setup used for the AMBench2018-01 benchmark of the NIST experiment [1].

Parameters	Value
Total number of layers	625
Average layer time	52 s
Layer height	20 μm
Contour scan speed	900 mm/s
Infill scan speed	800 mm/s
Hatch space	100 μm

are reported in Section 2.4.2. In addition, two build processes were performed using two different machines: the Additive Manufacturing Metrology Testbed (the NIST in-house build machine) and an EOS M270. The numerical simulations carried out in the present work and discussed in Section 2.4 refer to the process parameters of the EOS M270 machine, reported in Table 2.1.

After fabrication, part of the structure was removed from the build plate by electrical discharge machining, allowing the part to flex upward due to residual stresses. The eleven ridges created on the top surface of the structure (see Figure 2.1) were ground to obtain a smooth surface suitable for accurate displacement measurements.

Several measurements were performed at the NIST laboratories, including the upward displacement of the cantilever beam after partial removal from the build plate. The present work considers the residual strains and the upward displacement of the cantilever beam, both before and after removal from the build plate, providing the reference quantities for subsequent uncertainty quantification and model calibration.

2.4 Numerical modeling

The numerical model adopted to reproduce the NIST experiment described in Section 2.3 employs a part-scale, transient thermomechanical FE model of the AMBench2018-01 Inconel 625 cantilever beam. The geometry of the cantilever beam is shown in Figure 2.2. The modeling approach follows established part-scale practices for PBF-LB/M [15, 16, 46, 53–57]. The transient thermal analysis is advanced in time throughout the entire build, generating a layer-by-layer history of temperature fields that captures cumulative thermal exposure and cooling intervals. The temperature fields are then used as thermal loads for the quasi-static mechanical analysis, which is solved in a staggered manner. The weak (staggered) coupling between the thermal and mechanical analyses avoids a monolithic thermomechanical solution at each time step while retaining the dominant mechanisms responsible for thermal strains, plastic strains, and residual stress development.

All numerical simulations are carried out in *ANSYS 2021R2*, using built-in functionalities for element activation to represent layer addition, field transfer to map temperature histories to the structural analysis, and element deactivation to emulate the electrical-discharge machining cut performed prior to the evaluation of the residual-stress-induced upward displacement [15, 16, 53, 86].

The computational domain includes the build plate and the component, i.e., the cantilever beam. Powder is not explicitly represented at the part scale; instead, heat losses toward the powder and toward the environmental gas are modeled through distinct convection boundary conditions as discussed in Section 2.2. The energy input is not resolved at the melt-pool scale. The build sequence is imposed through a layer-by-layer activation strategy governed by measured or prescribed layer times and inter-layer dwell times. An activation temperature parameter T_A is introduced to trigger element activation and to approximate the onset of thermomechanical loading within each newly added layer.

Once the thermal analysis is completed, the mechanical analysis applies fixed constraints on the build plate during the build. The prescribed portion of the cantilever beam is then deactivated to reproduce the removal cut, enabling the evaluation of the upward displacement induced by the residual stress field. The quantities of interest include residual strains extracted along the line at $z = 11$ mm and displacements at selected measurement points, in accordance with Figure 2.2 and with the AMBench2018-01 benchmark [1, 2].

The remainder of the section describes in detail the modeling choices required to ensure predictive accuracy and reproducibility of the NIST AMBench2018-01 benchmark.

2.4.1 Meshing strategy

All numerical simulations employ a structured mesh composed of quadratic hexahedral finite elements. The use of a uniform mesh topology simplifies the layer-by-layer strategy and ensures consistent transfer of the transient temperature field to the quasi-static mechanical analysis. The build plate and the cantilever beam are meshed with different characteristic element sizes to

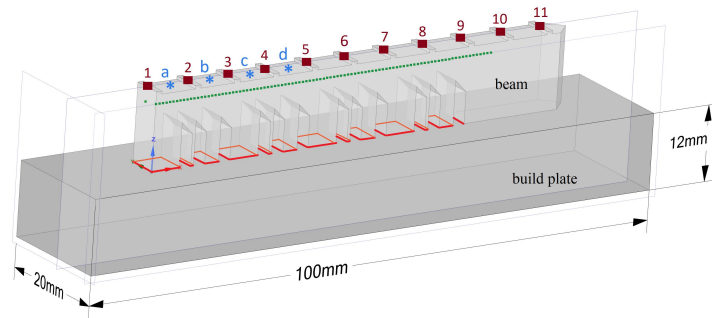


Figure 2.2: Part-scale model: Inconel 625 cantilever beam with length 75 mm, height 12 mm, and width 5 mm, printed on a build plate of dimensions 85 mm \times 12 mm \times 20 mm. Numbered (purple) and lettered (blue) points denote measurement locations used in the UQ analysis; the dotted (green) line at $z = 11$ mm identifies where residual strains are extracted; the solid (red) contour indicates the cut terminating at $x = 56$ mm.

reflect their distinct modeling roles. The cantilever beam requires finer spatial resolution to capture temperature gradients and stress redistribution during the process, whereas the build plate primarily acts to simulate heat losses during the thermal analysis and as a mechanical constraint during the mechanical analysis, and can therefore be meshed more coarsely.

In addition, in the present work, two mesh strategies are considered for the cantilever beam, referred to as the coarse mesh and the fine mesh. Representative meshes are shown in Figure 2.3. The coarse mesh adopts a characteristic element size of 0.5 mm, whereas the fine mesh refines the cantilever beam to 0.25 mm. In both cases, the mesh is aligned with the build direction, and element layers are stacked through the cantilever beam height to ensure a regular activation pattern during the thermal analysis. The build plate is meshed with a characteristic element size of 3 mm.

Table 2.2 reports the number of nodes, number of elements, number of finite element layers through the cantilever beam height (“nb. z -layers”), and indicative times required to complete the high-fidelity thermomechanical simulations. The z -layer count refers to the number of finite element layers across the 12 mm height of the cantilever beam and must not be confused with the physical process layers of the PBF-LB/M process, which are governed by the layer-by-layer strategy described in Section 2.4.3. As expected, the fine mesh increases the number of degrees of freedom by approximately one order of magnitude relative to the coarse mesh, resulting in a substantially increased computational cost. The choice of the two mesh strategies is related to the surrogate modeling strategy described in Section 2.5.

2.4.2 Material properties

The numerical simulations adopt temperature-dependent material properties of the Inconel 625 superalloy provided in the *ANSYS 2021R2* material library, in accordance with the AMBench2018-




Figure 2.3: Mesh strategies adopted for the part-scale PBF-LB/M thermomechanical model of the Inconel 625 cantilever beam. (a) Coarse mesh with characteristic element size 0.5 mm; (b) Fine mesh with characteristic element size 0.25 mm. The build plate is meshed with 3 mm elements in both cases.

Table 2.2: Mesh strategies adopted for the part-scale PBF-LB/M thermomechanical model of the Inconel 625 cantilever beam. For each mesh strategy, the table reports the number of nodes (nb. nodes), number of elements (nb. elements), number of finite element layers through the cantilever beam height (nb. z -layers), and the approximate time required to complete the high-fidelity thermomechanical numerical simulation.

mesh strategy	nb. nodes	nb. elements	nb. z -layers	sim. time (approx.)
coarse mesh	141095	30340	25	4 hours
fine mesh	1041027	240560	50	6 days

01 specification introduced in Section 2.3. The material is assumed homogeneous and isotropic, and all temperature-dependent quantities are supplied to the FE solvers as tabulated functions evaluated through piecewise-linear interpolation.

The thermo-elasto-plastic behaviour follows the constitutive framework introduced in Section 2.2. In this context, *ANSYS 2021R2* provides the temperature-dependent elastic constants and the bilinear isotropic hardening law used in the mechanical analysis. The reduction of stiffness and yield strength with increasing temperature, together with the associated tangent modulus, is illustrated in Figure 2.4. Mass density and melting temperature are treated as temperature-independent, taking the values $\rho = 8440 \text{ kg m}^{-3}$ and $T_{\text{melt}} = 1290 \text{ }^\circ\text{C}$.

The transient thermal analysis relies on the temperature-dependent specific heat, thermal conductivity, and coefficient of thermal expansion made available in the same material dataset. These properties govern heat storage, heat diffusion, and thermal-strain generation throughout

the build, and exhibit the typical monotonic trends observed in nickel-based superalloys over the temperature range relevant to PBF-LB/M.

The complete set of tabulated values used in the numerical simulations is reported in Tables 2.3 and 2.4, while Figure 2.5 summarises the temperature dependence of the main thermophysical and mechanical quantities. The thermal histories obtained in the part-scale numerical simulations remain within the tabulated range, thereby avoiding extrapolation. The adopted property set provides a reproducible and consolidated engineering description suitable for the uncertainty-quantification study in Section 2.5, where variability arises primarily from boundary heat-transfer coefficients and the activation temperature rather than from material behaviour.




Figure 2.4: Temperature-dependent bilinear isotropic hardening law for Inconel 625 adopted in the mechanical analysis (*ANSYS 2021R2*).

Table 2.3: Temperature-dependent elastic properties of Inconel 625 from *ANSYS 2021R2*: Young’s modulus and Poisson’s ratio.

Temperature [°C]	Young’s modulus [Pa]	Poisson’s ratio [-]
24	1.6E+11	0.278
538	1.2E+11	0.305
816	5.7E+10	0.33
982	3.8E+10	0.33
1093	2.1E+10	0.33

Table 2.4: Temperature-dependent thermophysical properties of Inconel 625 from *ANSYS 2021R2*: coefficient of thermal expansion, specific heat, and thermal conductivity.

Temperature [°C]	Thermal expansion [°C ⁻¹]	Specific heat [J kg ⁻¹ °C ⁻¹]	Thermal conductivity [W m ⁻¹ °C ⁻¹]
93	1.28E-05	427	10.8
204	1.31E-05	456	12.5
316	1.33E-05	481	14.1
427	1.37E-05	500	15.7
538	1.4E-05	536	17.5
649	1.48E-05	565	19.0
760	1.53E-05	590	20.8
871	1.58E-05	620	22.8
927	1.62E-05	645	25.2

2.4.3 Printing parameters

The part-scale numerical model does not resolve the laser path or the melt-pool scale. Each newly added layer is introduced by assigning a uniform activation temperature T_A , after which the layer cools according to the boundary conditions defined in Section 2.2. Within the numerical modeling framework, the printing parameters serve two functions: (i) they determine the temporal progression of the build by specifying the nominal duration associated with each physical layer; and (ii) they influence the temporal spacing between consecutive layer activations by defining the time interval between successive layer activations.

Consistent with the AMBench2018-01 benchmark [1] described in Section 2.3, a nominal duration of 52 s is assigned to each physical layer. Since the experiment involved four identical cantilever beams printed on the same build plate and only one cantilever beam is explicitly represented in the numerical model, a dwell time multiplier is introduced to reproduce the delay associated with sequential printing. A multiplier equal to 4 is adopted, reflecting the presence of the four nominally identical cantilever beams as shown in Figure 2.1. Geometric and scanning parameters follow the EOS M270 machine settings reported in Table 2.1.

Heat transfer through the boundaries is governed by the two convection coefficients introduced in Section 2.2.1: the powder convection coefficient h_p , applied to faces adjacent to the powder, and the gas convection coefficient h_g , applied to faces exposed to the environmental gas (see Section 2.2.1). Both coefficients enter the Neumann-type boundary conditions that define heat loss and are treated as uncertain input parameters in the UQ analyses (see Section 2.5), reflecting variability in powder properties, environmental conditions, and near-surface morphology. The environmental temperature and the build plate temperature complete the set of thermal boundary

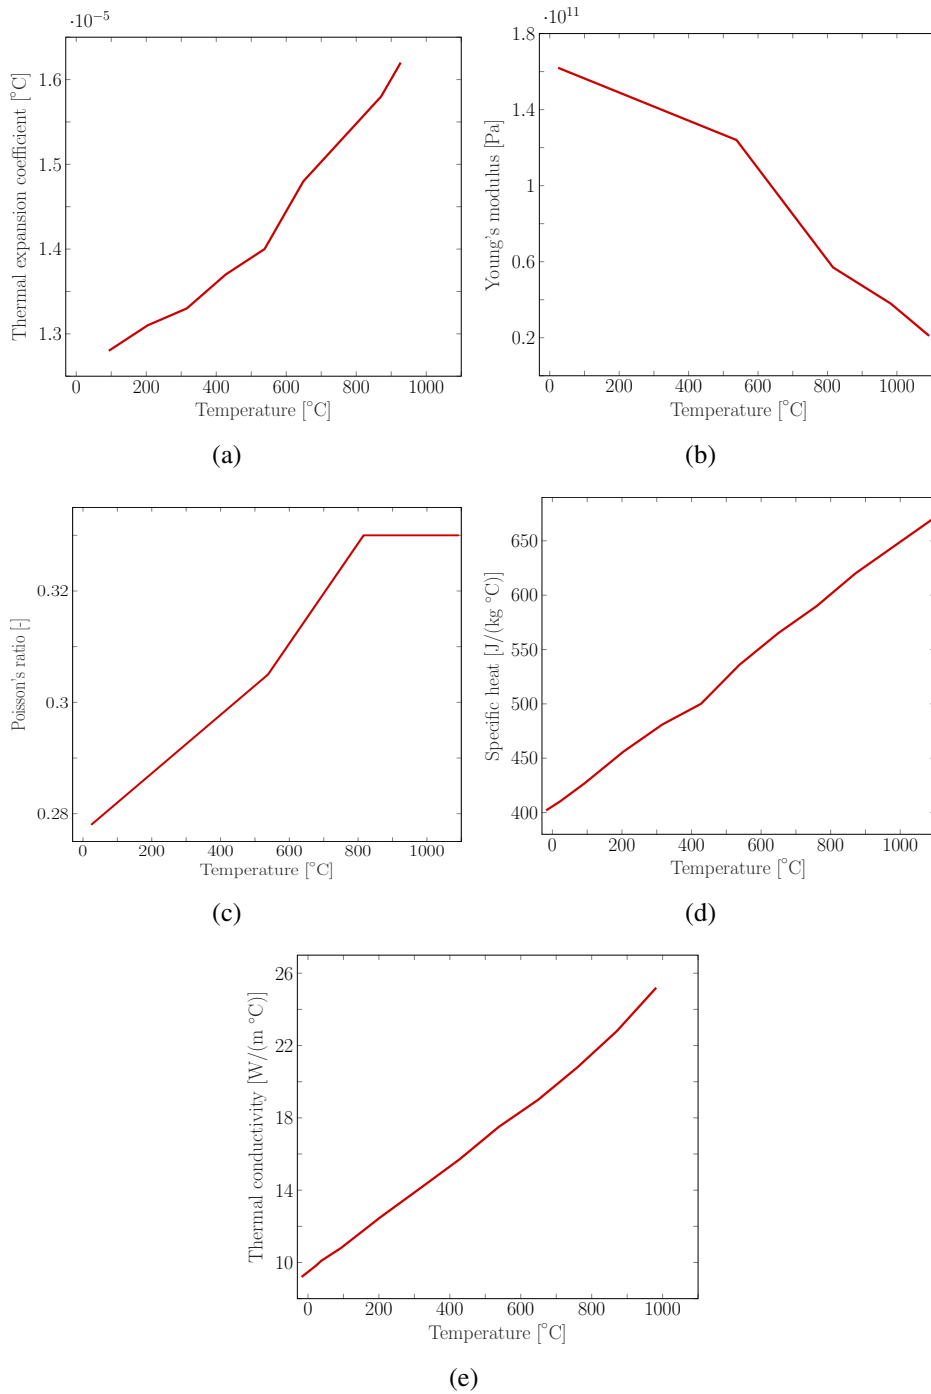


Figure 2.5: Temperature-dependent properties of Inconel 625 adopted in the simulations (*ANSYS 2021R2*): (a) coefficient of thermal expansion; (b) Young's modulus; (c) Poisson's ratio; (d) specific heat; (e) thermal conductivity.

inputs and are imposed as Dirichlet conditions during the build and cooling phases.

Table 2.5 summarizes the process parameters used to perform the numerical simulations.

Table 2.5: Machine and process parameters used in the part-scale PBF-LB/M numerical simulations of the Inconel 625 cantilever beam. Values reflect the AMBench2018-01 EOS M270 build, with the dwell time multiplier accounting for sequential printing of the four identical cantilever beams on the build plate.

Parameters	Value
deposition thickness	20 μm
hatch spacing	100 μm
dwell time multiplier	4
number of heat sources	1
contour scan speed	900 mm s^{-1}
infill scan speed	800 mm s^{-1}
laser power	100 W

2.4.4 Printing process

At the beginning of the numerical simulation, all elements belonging to the cantilever beam are deactivated and retained in a so-called “quiet” state, characterized by strongly reduced stiffness and conductivity. This guarantees numerical stability without contributing to either heat transfer or mechanical response. The build plate is active from the beginning of the analysis to represent the thermal and mechanical constraint imposed during the build.

The thermal analysis proceeds layer by layer in time. At each activation event, all elements belonging to the next super-layer are activated by switching them from the quiet state to the active state, and their nodal temperatures are instantaneously set to the activation temperature T_A . This operation idealizes the combined thermal effect of laser scanning within the physical layers. After activation, the newly added material cools according to the convection boundary conditions defined in Section 2.2. Throughout the transient thermal analysis, the top surface of the build plate is maintained at 80 °C and its lateral faces at 40 °C, consistent with the benchmark specification. All surfaces of the cantilever beam exchange heat with the powder and with the environmental gas through the effective convection coefficients h_p and h_g . Radiation effects are neglected, in accordance with standard part-scale modeling assumptions. The initial temperature of the entire domain is set to the environmental temperature of 20 °C, and after the final activation and dwell the system is cooled down to 20 °C to reproduce chamber equilibration before removal of the cantilever beam.

The removal of the cantilever beam from the build plate is reproduced by progressively deactivating the solid elements intersected by the cut shown in Figure 2.2. Deactivation releases the constraints along the cut interface, which subsequently becomes traction-free. The resulting stress redistribution induces an upward displacement of the cantilever beam, which is recorded at the designated measurement points for comparison with the AMBench2018-01 displacement field. Residual strains are extracted along the line at $z = 11$ mm, consistent with the AMBench2018 benchmark [1, 2].

2.5 Uncertainty quantification framework

The present section presents the uncertainty quantification (UQ) framework adopted to analyse the thermomechanical response of the part-scale PBF-LB/M model formulated in Sections 2.2 to 2.4. As outlined in Section 2.1, the methodology builds upon two complementary contributions previously developed by the author in [76, 77].

2.5.1 Uncertainty quantification based on single-fidelity sparse-grid surrogate modeling

This section introduces the uncertainty quantification (UQ) workflow based on a single-fidelity sparse-grid surrogate model, which enables the systematic execution of global sensitivity analysis (GSA), Bayesian inverse UQ, and data-informed forward UQ [76].

2.5.1.1 Sources of uncertainty

The parameters considered as uncertain in the present study are the base-10 logarithm of the powder convection coefficient, $\log h_p$, the base-10 logarithm of the gas convection coefficient, $\log h_g$ (see Section 2.2.1), and the activation temperature T_A (see Section 2.4.3). In UQ analyses, uncertain parameters are modeled as random variables.

In the absence of precise prior statistical information on the uncertain parameters, each parameter is modeled as an independent random variable with a uniform distribution over a prescribed interval. This choice corresponds to adopting a non-informative prior in the sense of the maximum entropy principle under bounded support.

The assumption of mutual independence reflects the lack of prior evidence of statistical correlation among the parameters and is consistent with standard practice in engineering UQ when joint statistical information is unavailable.

The intervals reported in Table 2.6 define the support of the prior distributions and are selected based on physically admissible ranges and values reported in the literature [87–91]. These bounds are intentionally chosen to be sufficiently wide to avoid artificially constraining the inference procedure, while ensuring physical consistency of the model inputs.

It is important to remark that the intervals defined in Table 2.6 play a dual role. On the one hand, they define the admissible parameter domain for the construction of the surrogate

model and the associated UQ analyses. On the other hand, they define the support of the prior distributions employed in the Bayesian inverse UQ framework. While the adoption of uniform priors reflects a non-informative modeling assumption, the bounds themselves are grounded in physical considerations and literature evidence. This distinction ensures consistency between the surrogate modeling stage and the subsequent Bayesian inference.

In this context, the prior distributions should be interpreted as weakly informative priors constrained to physically admissible ranges, rather than as fully data-driven probabilistic descriptions.

It is noted that, although h_g and h_p could have been treated directly as random variables, the logarithmic transformation proves more effective, as it enables spanning several orders of magnitude while assigning uniform resolution in the log-scale interval $[10^{-5}, 10^0]$. This parametrization ensures a balanced representation of both small and large values of the heat transfer coefficients.

The three uncertain parameters are collected in the vector

$$\mathbf{v} = (v_1, v_2, v_3) = (T_A, \log h_g, \log h_p)$$

and the following notation is introduced:

- $\Gamma_n = [a_n, b_n]$ denotes the interval associated with each uncertain parameter v_n , for $n = 1, 2, 3$;
- $\Gamma = \Gamma_1 \times \Gamma_2 \times \Gamma_3$ denotes the domain of \mathbf{v} , i.e. the hyper-rectangle $[a_1, b_1] \times [a_2, b_2] \times [a_3, b_3]$;
- $\rho_n(v_n)$ denotes the PDF of v_n . In accordance with the discussion above, the prior PDFs satisfy $\rho_{n,prior}(v_n) = \frac{1}{b_n - a_n}$, where the subscript ‘‘prior’’ indicates that the corresponding distribution incorporates only the a priori information specified in Table 2.6. After the inverse UQ procedure, these PDFs are updated to data-informed posterior PDFs, denoted by $\rho_{n,post}$;
- $\rho(\mathbf{v})$ denotes the joint PDF of the parameter vector. In particular, $\rho_{prior}(\mathbf{v})$ and $\rho_{post}(\mathbf{v})$ denote the joint prior and posterior PDFs, respectively. Under the assumption of mutual independence among the three parameters, $\rho_{prior}(\mathbf{v}) = \prod_{n=1}^3 \rho_{n,prior}(v_n)$. No analogous factorization can be assumed for ρ_{post} , since statistical dependence among the parameters may arise after the inverse UQ procedure;
- $u(\mathbf{x}, \mathbf{v})$ denotes the displacement along the z -direction at $\mathbf{x} \in \Omega$ corresponding to the parameter vector \mathbf{v} ;
- more generally, $\mathbf{f}(\mathbf{v}) : \Gamma \rightarrow \mathbb{R}^P$ denotes any QoI (output) of the numerical simulation, such as displacements or residual strains, emphasizing the dependence on \mathbf{v} . In the case of a scalar QoI ($P = 1$), the notation $f(\mathbf{v})$ is used.

Table 2.6: Uncertainty parameters of the PBF-LB/M printing process considered for the UQ analysis with associated interval.

uncertainty parameters	units	interval $[a_n, b_n]$
Activation temperature (T_A)	$^{\circ}C$	[1130; 1450]
log of gas convection coefficient ($\log h_g$)	–	[–5, 0]
log of powder convection coefficient ($\log h_p$)	–	[–5, 0]

2.5.1.2 Uncertainty quantification workplan

The fundamental premise of UQ is that, since \mathbf{v} is uncertain and described by a random vector equipped with an associated PDF, any output $\mathbf{f}(\mathbf{v})$ is likewise an uncertain quantity. The efficient characterization of the probability distribution of $\mathbf{f}(\mathbf{v})$ constitutes the final objective of the UQ analysis. In the present work, the QoI of interest is the field of residual strains along the cantilever beam.

In addition, experimental measurements are available for an auxiliary QoI, namely the displacement of the cantilever beam, which is incorporated into the analysis. The overall workflow combines forward and inverse UQ components and proceeds as follows.

1. **GSA** (see Section 2.5.1.4): the analysis investigates the contribution of each uncertain parameter v_n to the variability of the displacement field under the assumed prior distributions. The results indicate that the second parameter, i.e., the gas convection coefficient, exhibits a negligible influence on the displacements over the considered parameter domain and may therefore be fixed to a constant value in the subsequent analyses as a model reduction step.

It is noted that global sensitivity measures depend on both the assumed prior distributions and the structure of the numerical model. As such, the identification of negligible parameters should be interpreted as a global indication over the prescribed parameter domain and does not exclude the possibility of local sensitivity effects in restricted regions of the parameter space.

This reduction is introduced to improve computational tractability while preserving the dominant sources of variability identified at the global level. Accordingly, the updated prior PDF becomes $\rho_{prior,red}(\mathbf{v}) = \rho_1(v_1) \rho_3(v_3)$, and the reduced parameter space is $\Gamma_{red} = \Gamma_1 \times \Gamma_3$.

2. **Inverse UQ** (see Section 2.5.1.5): Bayesian inverse techniques are then employed to update the initial PDF $\rho_{prior,red}$ of the remaining uncertain parameters by incorporating the information provided by the displacement measurements. The outcome is a data-informed posterior PDF, ρ_{post} , which reflects the updated probabilistic description of the parameters conditioned on the available experimental data and is typically characterized by reduced dispersion relative to $\rho_{prior,red}$.

3. **Data-informed forward UQ** (see Section 2.5.1.6): the reduced parameter space Γ_{red} is sampled according to the posterior PDF ρ_{post} , and the corresponding residual strains are evaluated. These evaluations are then used to construct the probability distribution of the residual strains, thereby enabling a data-informed prediction of the QoI.

2.5.1.3 Sparse-grid surrogate modeling for uncertainty quantification

All the steps in the workplan require repeatedly solving the PBF-LB/M model for different values of \mathbf{v} to evaluate the corresponding QoIs $\mathbf{f}(\mathbf{v})$ (displacements, residual strains). Specifically, in *GSA* it is necessary to assess how variations in each parameter v_n affect the displacement field. In the *inverse UQ*, it is instead necessary to evaluate the compatibility between the model predictions and the available displacement data: parameter values \mathbf{v} that produce model responses that are far from the experimental observations are associated with low posterior probability. Finally, in *data-informed forward UQ*, evaluations of the residual strains are required to characterize their probability distribution.

To reduce the computational burden, the values of $\mathbf{f}(\mathbf{v})$ obtained from the evaluation of the PBF-LB/M model are replaced by suitable approximations constructed through surrogate modeling. Surrogate models are typically obtained through a two-step procedure. In the first step (“offline” or “training” phase), the PBF-LB/M model is evaluated for a limited set of carefully selected parameter values \mathbf{v} , yielding samples $\mathbf{f}(\mathbf{v}_1), \dots, \mathbf{f}(\mathbf{v}_M)$, from which an approximation of $\mathbf{f}(\mathbf{v})$ is constructed. In the second step (“online” or “evaluation” phase), the surrogate model is evaluated in place of the high-fidelity model, resulting in a significant reduction in computational cost.

In the present work, sparse-grid surrogate models are employed. These are established surrogate modeling techniques in uncertainty quantification, particularly suited for problems characterized by low-to-moderate dimensionality and sufficiently smooth model responses [92–94]. Alternative surrogate modeling approaches, such as Gaussian process regression, sparse polynomial chaos expansions, or neural-network-based methods, are often preferred in higher-dimensional or strongly nonlinear settings. However, for the present problem, characterized by a limited number of uncertain parameters and a sufficiently smooth dependence of the QoIs on the inputs, sparse-grid methods provide an appropriate and accurate approximation strategy.

In the present section, the basics of sparse grids are outlined, while additional details are provided in Appendix A.

Before introducing the definition of sparse-grid surrogate models, it is noted that the approximation of a QoI $\mathbf{f}(\mathbf{v})$ depends, among other factors, on the number of uncertain parameters N and on their associated PDFs. As a consequence, three different sparse-grid surrogate models are constructed: the first one for the *GSA* (three parameters, ρ_{prior}), the second one for the *inverse UQ* analysis (two parameters with $\rho_{prior,red}$), and the third one for the *data-informed forward UQ* analysis (two parameters with ρ_{post}). Table 2.7 summarizes the properties of these surrogate models.

The relatively small number of collocation points reflects the regularity of the considered QoIs with respect to the uncertain parameters. In the present application, the thermomechanical

Table 2.7: Properties of the three sparse-grid surrogate models used in this work.

sparse grid	used in	N	PDF	\mathcal{I}	w	univariate nodes	points	surrogate for
1	GSA	3	ρ_{prior}	\mathcal{I}_{max}	1	symm. Leja	27	11 z -displacements
2	inverse UQ	2	$\rho_{prior,red}$	\mathcal{I}_{sum}	3	symm. Leja	25	9 z -displacements
3	forward UQ	2	ρ_{post}	\mathcal{I}_{sum}	3	symm. Leja ($\log h_p$) symm. Gaussian Leja (T_A)	25	120 x -residual strains

response exhibits a sufficiently smooth dependence on the input parameters, which allows sparse-grid surrogate models to achieve accurate approximations with a limited number of model evaluations. This behavior is consistent with the underlying physics of the part-scale model, where the dominant effects are governed by diffusive thermal processes and gradually evolving mechanical fields.

In general, the sparse-grid surrogate model of a scalar-valued $f(\mathbf{v})$ for $\mathbf{v} \in \Gamma \subset \mathbb{R}^N$ with associated PDF $\rho(\mathbf{v})$ is denoted by $\mathcal{S}_{\mathcal{I}}f(\mathbf{v})$ and reads:

$$f(\mathbf{v}) \approx \mathcal{S}_{\mathcal{I}}f(\mathbf{v}) = \sum_{\mathbf{i} \in \mathcal{I}} c_{\mathbf{i}} \mathcal{U}_{\mathbf{i}}(\mathbf{v}), \quad c_{\mathbf{i}} := \sum_{\substack{\mathbf{j} \in \{0,1\}^N \\ \mathbf{i} + \mathbf{j} \in \mathcal{I}}} (-1)^{\|\mathbf{j}\|_1}.$$

- $\mathbf{i} = [i_1, i_2, \dots, i_N] \in \mathbb{N}_+^N$ is a multi-index with $i_n \geq 1$;
- $\mathcal{I} \subset \mathbb{N}_+^N$ is a downward-closed multi-index set (see Appendix A);
- $\mathcal{U}_{\mathbf{i}}(\mathbf{v})$ is a tensor Lagrangian interpolant constructed on a Cartesian grid with

$$(2i_1 - 1) \times \dots \times (2i_N - 1) \tag{2.14}$$

points;

- $c_{\mathbf{i}}$ are the combination technique coefficients.

The sparse-grid surrogate model is thus obtained as a linear combination of tensor interpolants defined over different grids. The choice of the multi-index set \mathcal{I} is therefore crucial.

A simple choice is

$$\mathcal{I}_{max} = \{\mathbf{i} \in \mathbb{N}_+^N : \max_n (i_n - 1) \leq w\},$$

which leads to a full tensor grid with $(2w + 1)^N$ points and is therefore computationally prohibitive for moderate N .

A more efficient alternative is

$$\mathcal{I}_{sum} = \{\mathbf{i} \in \mathbb{N}_+^N : \sum_{n=1}^N (i_n - 1) \leq w\},$$

which yields a sparse sampling of the parameter space while retaining good approximation properties.

The union of all points required by the interpolants $\mathcal{U}_i(\mathbf{v})$ defines the sparse grid. Figure 2.6 illustrates an example.

The selection of univariate collocation points is also critical. Equispaced nodes are avoided due to the Runge phenomenon [95]. Instead, non-equispaced nodes consistent with the underlying PDFs are used.

For uniform distributions (ρ_{prior} and $\rho_{prior,red}$), symmetric Leja points are employed, while for the posterior-based surrogate (ρ_{post}), symmetric Gaussian Leja points are used for T_A and symmetric Leja points for $\log h_p$. These choices ensure good interpolation properties and efficient sampling of the parameter space.

Finally, the extension to vector-valued QoIs is straightforward, as the same sparse grid can be used for each component.

Having introduced the surrogate modeling framework, the detailed description of each step of the UQ workflow can now be provided.

2.5.1.4 Global sensitivity analysis

Global sensitivity analysis (GSA) is employed to quantify the influence of the uncertain parameters (see Table 2.6) on the variability of the displacement field. In the present work, the Sobol decomposition method [96–100] is employed.

The Sobol method is based on a variance decomposition analogous to the classical ANOVA expansion and enables the partition of the total variance of the QoIs into contributions associated with each input parameter and their interactions. Two sets of sensitivity measures are considered, namely the principal Sobol indices and the total Sobol indices. The former quantify the individual contribution of each parameter to the variance of the QoIs, whereas the latter account for the total contribution of each parameter, including all interaction effects with the remaining inputs.

In practice, the computation of Sobol indices requires repeated evaluations of the QoIs over the parameter space. To alleviate the associated computational cost, the analysis is performed by replacing the high-fidelity PBF-LB/M model with the sparse-grid surrogate model introduced in Section 2.5.1. This surrogate-based approach enables the estimation of the sensitivity indices while preserving the accuracy of the variance decomposition.

It is important to note that polynomial chaos expansions (PCE) combined with ANOVA-type decompositions (often referred to as PCE–ANOVA or PCE–ANCOVA) represent a widely adopted and highly efficient framework for global sensitivity analysis, particularly in the context of surrogate modeling [101–104]. In such approaches, Sobol indices can be obtained directly from the coefficients of the expansion, leading to significant computational advantages. These approaches are particularly effective when an explicit polynomial representation of the model response is available.

However, in the present work, sparse-grid surrogate models are employed in a non-intrusive interpolation setting, which provides a flexible approximation framework compatible with the structure of the underlying high-fidelity model and the adopted UQ workflow. This choice is consistent with the low-dimensional setting of the problem and with the smooth dependence of the QoIs on the uncertain parameters. The use of sparse grids allows for an accurate approximation

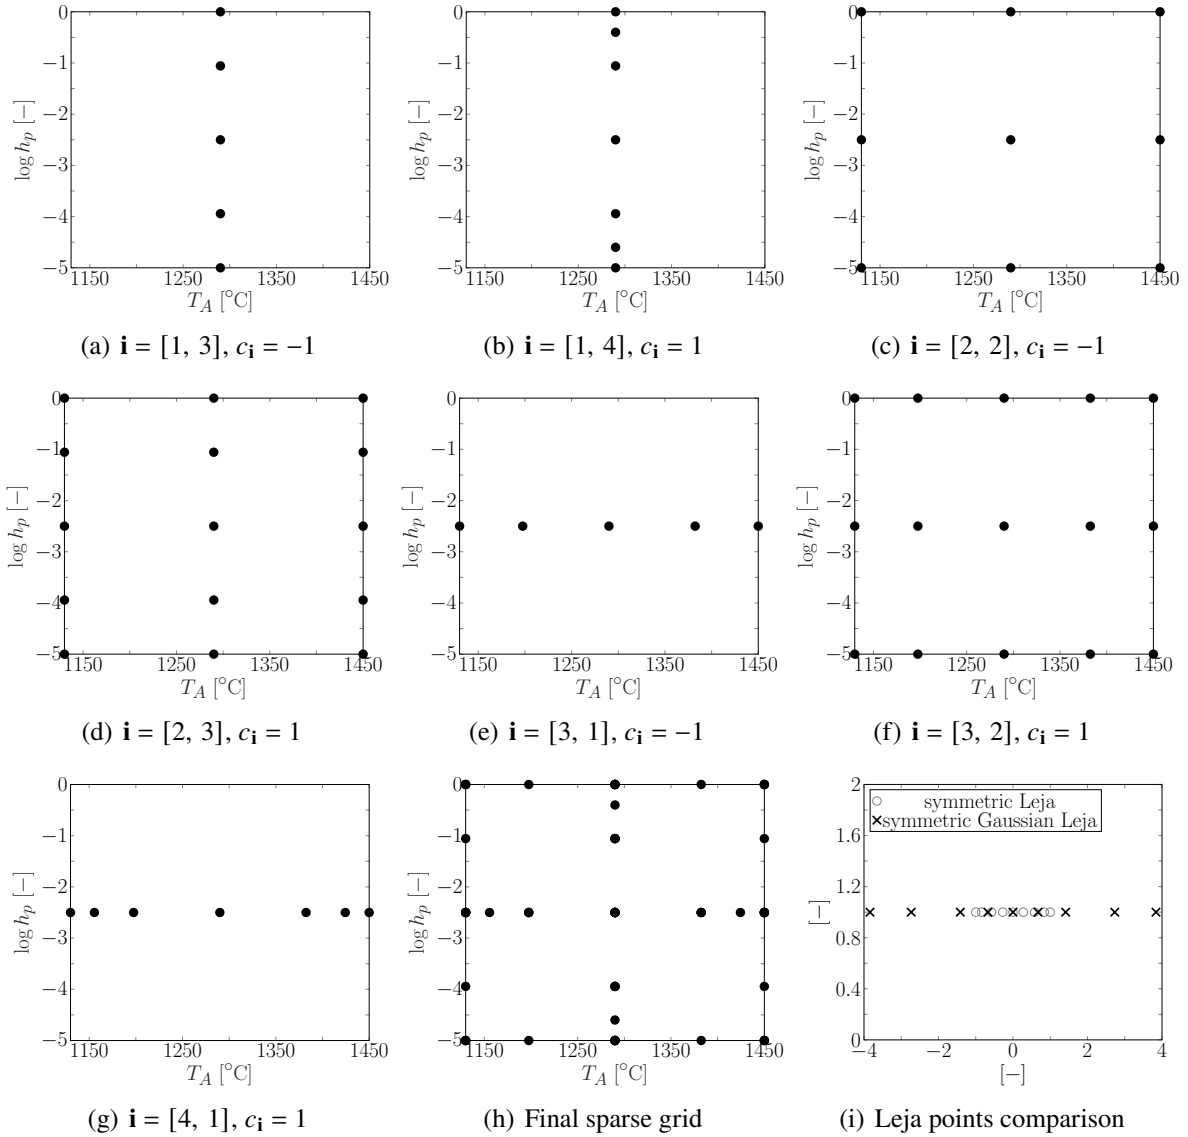


Figure 2.6: Breakdown of the sparse grid (h) into the tensor grids composing it (a)–(g). Figure (i) illustrates the different families of collocation points used in this work: symmetric Leja points for a uniform PDF on $\Gamma_n = [-1, 1]$, and symmetric Gaussian Leja points for a Gaussian PDF with zero mean and unit standard deviation.

of the QoIs with a limited number of model evaluations, enabling the subsequent computation of Sobol indices without requiring an explicit polynomial expansion of the model response.

The resulting sensitivity indices provide a global assessment of the relative importance of the uncertain parameters over the prescribed parameter domain and form the basis for the model reduction strategy described in Section 2.5.1.

2.5.1.5 Inverse uncertainty quantification analysis

A Bayesian inverse approach [60–62] is adopted to perform the inverse UQ analysis. More specifically, Bayesian inversion consists in updating the PDF of the two remaining uncertain parameters T_A and $\log h_p$ from the prior distribution $\rho_{prior,red}$ to a posterior PDF ρ_{post} , which incorporates the information provided by a set of displacement measurements at K positions $\mathbf{x}_{1,meas}, \dots, \mathbf{x}_{K,meas}$ (see Section 2.5.1.7.2.1 for details on the locations of $\mathbf{x}_{k,meas}$).

Ideally, experimental displacement measurements would be employed in the inverse UQ analysis. However, as a preliminary step, the present work considers a set of imperfect (noisy) synthetic data \tilde{u} , obtained by evaluating the part-scale thermomechanical model at a prescribed parameter vector $\bar{\mathbf{v}}$ (referred to as the *target values*), and subsequently perturbing the resulting displacement field by independent Gaussian noise realizations to mimic measurement errors:

$$\begin{cases} \tilde{u}_k = u(\mathbf{x}_{k,meas}, \bar{\mathbf{v}}) + \varepsilon_k, & k = 1, \dots, K, \\ \varepsilon_k \sim \mathcal{N}(0, \sigma_{\varepsilon_k}^2). \end{cases} \quad (2.15)$$

This setting enables a controlled assessment of the inversion procedure, isolating methodological aspects from additional discrepancies that could arise due to model-form errors. For simplicity, the noise level is assumed constant, i.e., $\sigma_{\varepsilon_k} = \bar{\sigma}$ for all k .

To formulate the Bayesian inverse problem, the misfits $M_k(\mathbf{v})$ between the synthetic data and the model-predicted displacements are introduced:

$$M_k(\mathbf{v}) := \tilde{u}_k - u(\mathbf{x}_{k,meas}, \mathbf{v}), \quad k = 1, \dots, K. \quad (2.16)$$

The posterior PDF ρ_{post} is obtained through Bayes' theorem [60, 61]:

$$\rho_{post}(\mathbf{v}) = \mathcal{L}(\mathbf{v} | \tilde{u}_k) \rho_{prior,red}(\mathbf{v}) \frac{1}{C}, \quad (2.17)$$

where C is a normalization constant and $\mathcal{L}(\mathbf{v} | \tilde{u}_k)$ denotes the likelihood function, which quantifies the plausibility of the parameter vector \mathbf{v} given the observed data. Under the assumption of independent Gaussian measurement noise (see Equation (2.15)), the likelihood reads:

$$\mathcal{L}(\mathbf{v} | \tilde{u}_k) = \prod_{k=1}^K \frac{1}{\sqrt{2\pi\sigma_{\varepsilon_k}^2}} \exp\left(-\frac{M_k^2(\mathbf{v})}{2\sigma_{\varepsilon_k}^2}\right). \quad (2.18)$$

In practice, the repeated evaluations of the model response $u(\mathbf{x}_{k,meas}, \mathbf{v})$ required for the construction of the likelihood would be computationally prohibitive if performed using the high-fidelity PBF-LB/M model. To overcome this limitation, the likelihood function is evaluated by replacing the high-fidelity model with the sparse-grid surrogate model introduced in Section 2.5.1. This surrogate-based formulation ensures computational tractability while preserving an accurate representation of the dependence of the QoIs on the uncertain parameters.

The final objective of the UQ workflow (step 3) is to sample the reduced parameter space Γ_{red} according to ρ_{post} , in order to propagate the updated uncertainty to the residual strains (data-informed forward UQ). Sampling the posterior distribution can be performed using Markov Chain Monte Carlo (MCMC) methods [105, 106], which may nevertheless remain computationally intensive even in the presence of surrogate models.

A significant reduction in computational cost is achieved by approximating ρ_{post} with a Gaussian distribution [62], characterized by an appropriate mean vector and covariance matrix. This approximation corresponds to a Laplace approximation of the posterior distribution around its maximum a posteriori (MAP) estimate [107, 108]. The validity of this approximation is discussed in the following subsection, which details the computation of the mean and covariance matrix of the Gaussian approximation of ρ_{post} . In this context, the second sparse-grid surrogate model reported in Table 2.7 plays a central role.

2.5.1.5.1 Gaussian approximation of posterior probability density function

The Gaussian approximation introduced in the previous subsection corresponds to a Laplace approximation of the posterior PDF around its maximum. Accordingly, the mean of the Gaussian approximation is taken as the maximum of the posterior PDF, \mathbf{v}_{MAP} (MAP), defined as:

$$\mathbf{v}_{MAP} := \arg \max_{\mathbf{v} \in \Gamma_{red}} \rho_{post}(\mathbf{v}) = \arg \max_{\mathbf{v} \in \Gamma_{red}} \mathcal{L}(\mathbf{v} | \tilde{u}_k), \quad (2.19)$$

where the second equality follows from the fact that $\rho_{prior,red}$ and C are constant with respect to \mathbf{v} .

This maximization is equivalent to a nonlinear least-squares problem, obtained by minimizing the negative log-likelihood:

$$\mathbf{v}_{MAP} = \arg \min_{\mathbf{v} \in \Gamma_{red}} [-\log \mathcal{L}(\mathbf{v} | \tilde{u}_k)] = \arg \min_{\mathbf{v} \in \Gamma_{red}} LS(\mathbf{v}), \quad (2.20)$$

$$LS(\mathbf{v}) = \sum_{k=1}^K M_k^2(\mathbf{v}). \quad (2.21)$$

The functional $[-\log \mathcal{L}(\mathbf{v} | \tilde{u}_k)]$ is commonly referred to as the *negative log-likelihood*. The minimization of $LS(\mathbf{v})$ requires multiple evaluations of the model response $u(\mathbf{x}_{k,meas}, \mathbf{v})$.

To reduce the computational cost, the high-fidelity model is replaced by the sparse-grid surrogate model reported in the second row of Table 2.7, leading to the modified misfit definition:

$$M_k(\mathbf{v}) := \tilde{u}_k - \mathcal{S}_{I_{sum}} u(\mathbf{x}_{k,meas}, \mathbf{v}), \quad k = 1, \dots, K.$$

The use of the surrogate model makes the evaluation of the likelihood computationally inexpensive and, in principle, enables the use of sampling-based approaches such as MCMC for posterior exploration. However, in the present work, a Laplace approximation is adopted in order to obtain an explicit Gaussian approximation of the posterior distribution, which provides a compact representation of parameter uncertainty in terms of mean and covariance and enables efficient propagation in the subsequent forward UQ step.

Once \mathbf{v}_{MAP} has been computed, the covariance matrix of the Gaussian approximation of ρ_{post} is obtained as:

$$\Sigma_{post} = \bar{\sigma}_{MAP}^2 \left(J_u^T J_u + \sum_{k=1}^K M_k H_{\tilde{u}_k} \right)^{-1}, \quad (2.22)$$

which corresponds to the inverse of the Hessian of the negative log-posterior evaluated at \mathbf{v}_{MAP} , consistently with the Laplace approximation framework [36, 37].

The quantities appearing in Equation (2.22) are defined as follows:

- $\bar{\sigma}_{MAP}^2$ is an estimate of the noise variance, obtained via the standard plug-in estimator:

$$\bar{\sigma}^2 \approx \bar{\sigma}_{MAP}^2 = \frac{1}{K} \sum_{k=1}^K (\tilde{u}_k - u(\mathbf{x}_{k,meas}, \mathbf{v}_{MAP}))^2 = \frac{1}{K} LS(\mathbf{v}_{MAP}); \quad (2.23)$$

- J_u is the Jacobian matrix of the displacement field with respect to \mathbf{v} , evaluated at \mathbf{v}_{MAP} ;
- $H_{\tilde{u}_k}$ denotes the Hessian of $u(\mathbf{x}_{k,meas}, \mathbf{v})$ with respect to \mathbf{v} , evaluated at \mathbf{v}_{MAP} .

The evaluation of the Jacobian and Hessian matrices is performed using finite-difference approximations based on the sparse-grid surrogate model, which enables efficient computation of first- and second-order derivatives with negligible additional cost.

The resulting Gaussian approximation provides a local characterization of the posterior distribution around the MAP estimate. The validity of this approximation depends on the degree of nonlinearity of the forward model and on the concentration of the posterior, and is discussed in the following section.

2.5.1.6 Data-informed forward uncertainty quantification analysis

Efficient characterization of the uncertainty in $\mathbf{f}(\mathbf{v})$ constitutes the main objective of the data-informed forward UQ analysis. This step relies on the posterior PDF ρ_{post} of the parameters in order to quantify the uncertainty in the prediction of the residual strains of the cantilever beam, given the uncertainty in T_A and $\log h_p$ as described by ρ_{post} .

More precisely, the aim is to approximate the PDF of the residual strains at L locations $\mathbf{x}_{j,str}$ along the x -direction, i.e. $\varepsilon_{xx}(\mathbf{x}_{j,str}, \mathbf{v})$ (see Section 2.5.2.3.1 for details on the locations of $\mathbf{x}_{j,str}$).

To this end, a sparse-grid surrogate model is first constructed for the residual strains, denoted by $\mathcal{S}_{I_{sum}} \varepsilon_{xx}(\mathbf{x}_{j,str}, \mathbf{v})$, at each of the L locations, using the specifications reported in the third row of Table 2.7. The accuracy of these surrogate models with respect to the high-fidelity model evaluations $\varepsilon_{xx}(\mathbf{x}_{j,str}, \mathbf{v})$ is verified a posteriori and found to be sufficient for the intended analysis (see Section 2.5.1.7.3).

The propagation of uncertainty is then performed by sampling the posterior distribution ρ_{post} . In the present work, 10^4 samples of \mathbf{v} are generated and used to evaluate the surrogate model, yielding corresponding realizations of the residual strains at each location.

It is noted that, due to the negligible computational cost of evaluating the sparse-grid surrogate model, significantly larger sample sizes (e.g. 10^5 – 10^6) could be employed without additional

computational burden. However, numerical experiments indicate that 10^4 samples are sufficient to ensure convergence of the estimated statistics and of the resulting probability density functions, with no appreciable changes observed for larger sample sizes.

Finally, the PDF of the residual strain at each location is approximated by applying a kernel density estimation method [109, 110] to the surrogate-based samples.

2.5.1.7 Results and discussion

This section reports the numerical results obtained from the UQ workflow described above, including the outcomes of the global sensitivity analysis (GSA), the inverse UQ analysis, and the data-informed forward UQ analysis.

All high-fidelity simulations of the PBF-LB/M process are performed on an HPC server equipped with 128 AMD EPYC 7702 @ 1.67 GHz CPU cores and 376 GB of RAM. The UQ analyses are implemented in MATLAB, using the *Sparse-Grids Matlab-Kit* [80].

2.5.1.7.1 Global sensitivity analysis

The global sensitivity analysis described in Section 2.5.1.4 is performed by considering as QoI the displacements at the centers of the eleven ridges of the cantilever beam (see Figure 2.2), i.e. a vector-valued QoI with components $\mathbf{f}_j(\mathbf{v}) = u(\mathbf{x}_{j,GSA}, \mathbf{v})$ for $j = 1, \dots, 11$. This yields 11 sets of Sobol indices, reported in Figure 2.7.

The Sobol indices are evaluated using the sparse-grid surrogate model introduced in Section 2.5.1, which replaces the high-fidelity PBF-LB/M model and enables an efficient computation of the variance-based sensitivity measures.

The sets corresponding to the first eight ridges, located further from the end of the removal area (Figure 2.2), exhibit consistent trends. In contrast, the remaining three sets are considered unreliable due to their proximity to the removal area and are therefore excluded from the subsequent analysis. In these regions, the vertical displacements are essentially negligible for all parameter values, which leads to an ill-conditioned variance decomposition and consequently to unreliable Sobol indices.

The results indicate that the parameters with the largest principal and total Sobol indices, and therefore the greatest influence on the displacements, are T_A and $\log h_p$, while $\log h_g$ has a negligible effect.

This observation motivates a reduction of the parameter space, whereby only T_A and $\log h_p$ are retained as uncertain parameters in the subsequent analyses, while $\log h_g$ is fixed to a constant value.

Specifically, $\log h_g$ is set to -5 , corresponding to the lower bound of the prior interval reported in Table 2.6. This choice is consistent with the Sobol analysis, which indicates negligible sensitivity of the QoI with respect to this parameter over the entire prior domain. Selecting the lower bound provides a conservative and physically admissible value within the prescribed range, corresponding to a regime in which heat losses through the surrounding gas are minimal compared to the contribution of powder convection.

Fixing $\log h_g$ in this manner reduces the dimensionality of the parameter space while preserving the dominant sources of variability identified by the global sensitivity analysis.

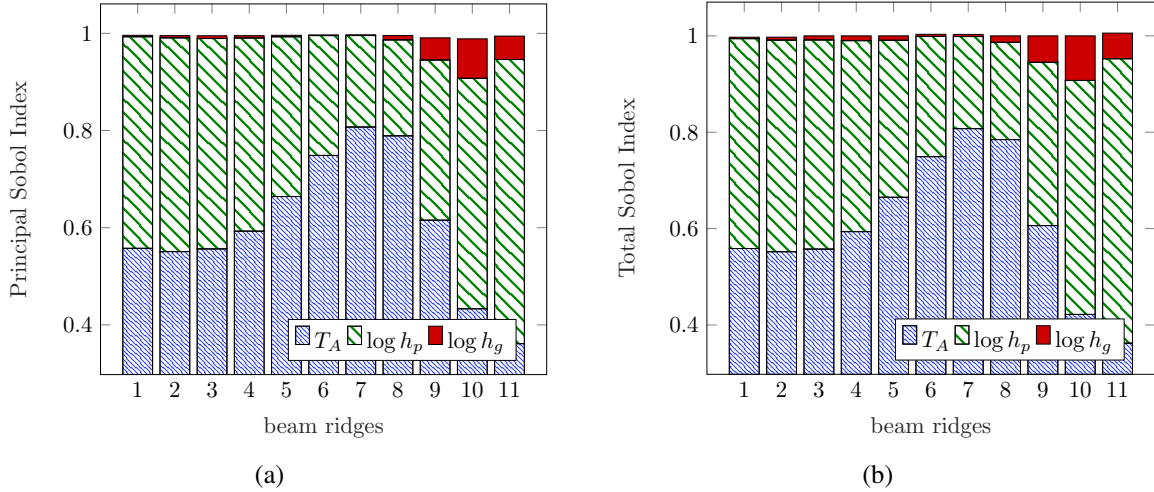


Figure 2.7: Magnitude of Sobol indices based on the displacements evaluated at the 11 ridges of the cantilever beam (QoI) for the three-parameter model: (a) Principal Sobol Indices that measure the *individual* contribution of each parameter to the variance of the QoI; (b) Total Sobol Indices that quantify the contributions of each parameter *combined with the others*.

2.5.1.7.2 Inverse uncertainty quantification

This section presents the results of the Bayesian inverse uncertainty quantification analysis described in Section 2.5.1.5. In particular, it reports the computation of the maximum a posteriori (MAP) estimate [107, 108], the Gaussian approximation of the posterior distribution obtained via the Laplace approximation [111, 112], and the corresponding reduction of uncertainty with respect to the prior distributions.

2.5.1.7.2.1 Surrogate modeling

A new sparse-grid surrogate model depending only on T_A and $\log h_p$ is constructed for use within the inverse UQ analysis. This surrogate is defined in the reduced parameter space identified by the GSA and is therefore consistent with the reduced prior distribution $\rho_{prior,red}$ and the subsequent Bayesian inversion. Constructing a surrogate directly in the reduced space improves both efficiency and approximation accuracy by avoiding unnecessary sampling along directions identified as non-influential.

As already mentioned, the new sparse grid employs the index set \mathcal{I}_{sum} and is based on 25 evaluations of the PBF-LB/M model (cf. row 2 of Table 2.7), corresponding to the points reported in Figure 2.8(a).

A vector-valued QoI is considered, with components $\mathbf{f}_k(\mathbf{v}) = u(\mathbf{x}_{k,meas}, \mathbf{v})$ for $k = 1, \dots, 9$, i.e., 9 measurement locations $\mathbf{x}_{k,meas}$ are selected at the first 5 ridges and at the 4 midpoints

between the respective ridges up to $x = 28.5$ mm (see Figure 2.2). This choice ensures that the measurement locations remain sufficiently distant from the end of the removal area ($x = 56$ mm), where numerical instabilities may affect the solution.

The corresponding 9 surrogate models exhibit similar behavior; therefore, in the following, results for the first ridge are reported. Figure 2.8(b) shows the surrogate model $\mathcal{S}_{I_{sum}} u(\mathbf{x}_{1,meas}, \mathbf{v})$, which exhibits a monotonic dependence on both parameters.

Before proceeding with the inversion, the accuracy of the sparse-grid surrogate model is assessed through a convergence study. Specifically, $M = 50$ random parameter pairs $\mathbf{v} = (T_A, \log h_p)$ are generated and, for each of them, the displacements at the 9 measurement locations are computed using the high-fidelity model. The corresponding surrogate predictions are then evaluated for increasing sparse-grid levels $w = 0, 1, 2, 3$.

It is emphasized that these $M = 50$ samples are used exclusively as an independent test set for error estimation and are not involved in the construction of the surrogate model.

The following error metrics are considered: the pointwise prediction error E_{PPE} and the quantity denoted as E_{MSE} , which, despite the notation, represents a root mean square *relative* error obtained by normalizing the squared residuals with respect to the corresponding displacement magnitude:

$$E_{PPE} = \max_{i=1,\dots,M} \frac{|u(\mathbf{x}_{k,meas}, \mathbf{v}_i) - \mathcal{S}_{I_{sum}} u(\mathbf{x}_{k,meas}, \mathbf{v}_i)|}{|u(\mathbf{x}_{k,meas}, \mathbf{v}_i)|}, \quad (2.24)$$

$$E_{MSE} = \sqrt{\frac{1}{M} \sum_{i=1}^M \frac{(u(\mathbf{x}_{k,meas}, \mathbf{v}_i) - \mathcal{S}_{I_{sum}} u(\mathbf{x}_{k,meas}, \mathbf{v}_i))^2}{(u(\mathbf{x}_{k,meas}, \mathbf{v}_i))^2}}. \quad (2.25)$$

The normalization by the magnitude of the reference solution provides a pointwise relative error measure that is directly interpretable in terms of percentage deviation. It is acknowledged that alternative normalization strategies, such as scaling with respect to the standard deviation of the sample, may offer improved robustness, particularly in situations where the reference quantity approaches zero. In the present case, however, the displacement values at the selected measurement locations are strictly non-zero and sufficiently bounded away from zero, so that the adopted normalization does not introduce numerical instabilities and remains representative of the approximation error.

The results of the convergence test reported in Figures 2.9(a) and 2.9(b) show that the sparse-grid surrogate model with $w = 3$ (25 sparse-grid points) provides a sufficiently accurate approximation for the inverse UQ analysis. In particular, the maximum relative error is approximately 1%, while the root mean square relative error is even smaller.

The accuracy of the surrogate model is further illustrated in Figure 2.10, where the displacements obtained from the high-fidelity model are compared with those predicted by the surrogate model. As the sparse-grid level increases, the surrogate predictions progressively align with the high-fidelity results, confirming the reliability of the surrogate model for the subsequent inversion step.

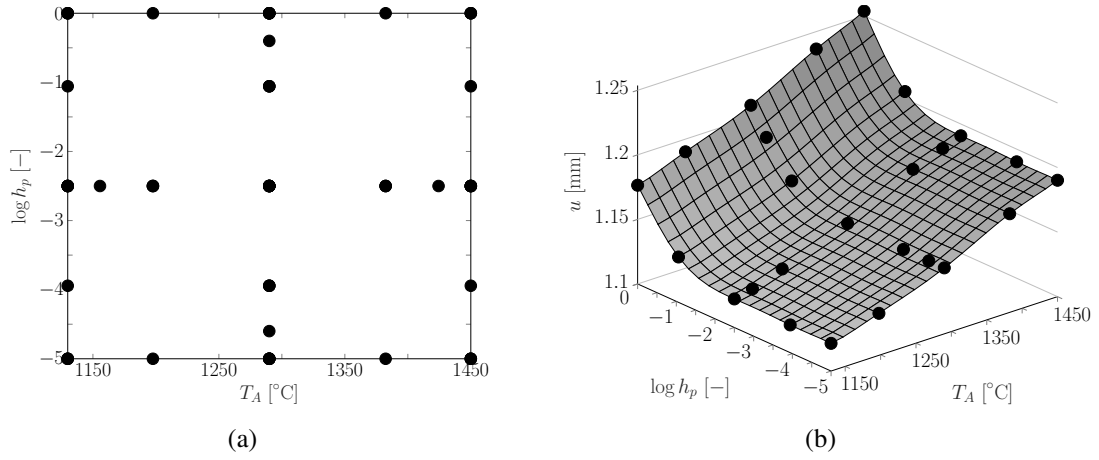


Figure 2.8: Sparse-grid surrogate model construction for the inverse UQ analysis with $w = 3$ (25 sparse-grid points) for the first ridge of the cantilever beam (i.e., $\mathbf{x}_{1,meas} = (0.5, 2.5, 12.5)$): (a) Sparse grid; (b) Surrogate model.




Figure 2.9: Results of the sparse-grid surrogate model for inverse UQ analysis. Convergence test for the first ridge of the cantilever beam (i.e., $\mathbf{x}_{1,meas} = (0.5, 2.5, 12.5)$): (a) Pointwise prediction error E_{PPE} ; (b) Root mean square error E_{MSE} (defined as relative error, see text).

2.5.1.7.2.2 Bayesian inverse analysis

As discussed in Section 2.5.1.5, the Bayesian inversion is performed using synthetic noisy displacement data \tilde{u}_k , $k = 1, \dots, 9$, generated according to Equation (2.15). The noise level is set to $\bar{\sigma} = 10^{-2}$ and the target parameter vector is chosen as $\bar{\mathbf{v}} = (T_A; \log h_p) = (1339.8 \text{ °C}; -3.75)$.

The synthetic displacement data are reported in Figure 2.11, together with the corresponding displacements obtained from the high-fidelity thermomechanical model at the target value $\bar{\mathbf{v}}$. The

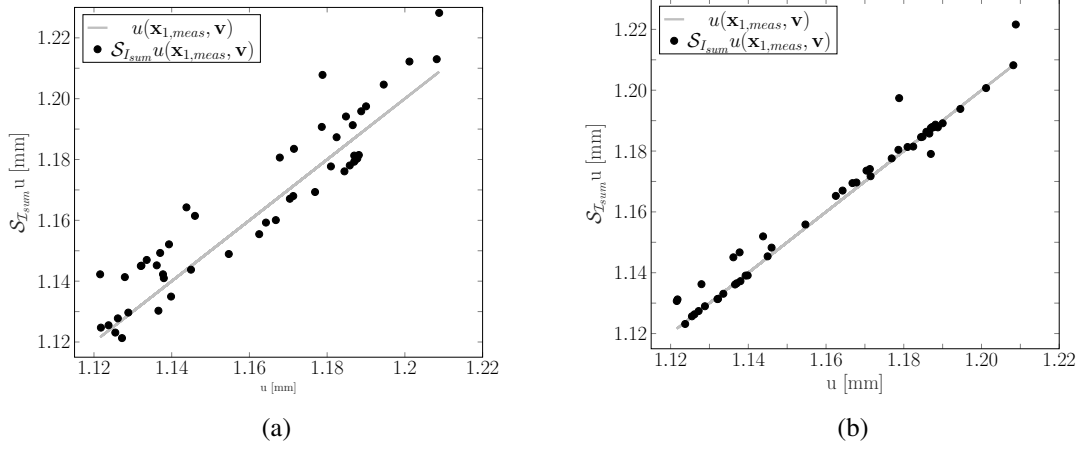


Figure 2.10: Results of the sparse-grid surrogate model for inverse UQ analysis. Comparison between sparse-grid surrogate model displacements and part-scale thermomechanical model displacements for the first ridge of the cantilever beam (i.e., $\mathbf{x}_{1,meas} = (0.5, 2.5, 12.5)$): (a) Sparse grid with level $w = 1$; (b) Sparse grid with level $w = 3$.

figure also displays the associated error bars. Unlike standard representations, the error bars are centered at the exact (yet unknown) displacement values $u(\mathbf{x}_{k,meas}, \bar{\mathbf{v}})$ and represent the interval $u(\mathbf{x}_{k,meas}, \bar{\mathbf{v}}) \pm 3\bar{\sigma}$, which contains approximately 99% of the realizations of the Gaussian noise model.

The measurement locations $\mathbf{x}_{k,meas}$ coincide with those introduced in the previous subsection. The adopted dataset is sufficient for the present analysis; additional tests performed with larger datasets do not significantly alter the results.

The mean of the Gaussian approximation of ρ_{post} , i.e. \mathbf{v}_{MAP} (see Equation (2.20)), is computed by minimizing the least-squares functional $LS(\mathbf{v})$. This minimization is carried out using the derivative-free Nelder–Mead algorithm (implemented in Matlab through the `fminsearch` routine). To improve robustness, the optimization is repeated from multiple initial guesses.

Three local minima are identified (see Figure 2.12(b)), one of which lies outside the admissible range of T_A (see Table 2.6) and is therefore discarded. The two remaining minima exhibit very similar values of $T_A \approx 1341$ °C, but significantly different values of $\log h_p$ (approximately -5 and -3).

This behavior indicates a lack of identifiability of the parameter $\log h_p$ based on the available displacement data. In such situations, the Laplace approximation may yield a Gaussian posterior that is artificially spread along directions of weak sensitivity, potentially extending beyond the support of the prior distribution.

For this reason, a Gaussian approximation is retained only for the parameter T_A , while $\log h_p$ is modeled through a uniform posterior distribution over the interval where the likelihood remains approximately constant.




Figure 2.11: Displacements $u(\mathbf{x}_{k,meas}, \bar{\mathbf{v}})$ obtained from part-scale thermomechanical analysis (gray curve) and synthetic displacements data \tilde{u}_k (black marker). Error bars associated with the measurements are also reported. Unlike standard error bar plots however, the error bars are centered at the exact (yet unknown) values on the displacements $u(\mathbf{x}_{k,meas}, \bar{\mathbf{v}})$, and show the ranges within which the actual noisy measurements are most likely found, according to Equation (2.15), i.e., $u(\mathbf{x}_{k,meas}, \bar{\mathbf{v}}) \pm 3\bar{\sigma}$.

This interpretation is supported by the inspection of the isolines of the least-squares functional (Figure 2.12(b)), which reveal an elongated flat region along the $\log h_p$ direction. Further insight is obtained by analyzing one-dimensional profiles of LS : fixing $\log h_p = -3$ yields a parabolic profile in T_A (Figure 2.12(c)), whereas fixing $T_A = 1341^\circ\text{C}$ leads to an almost flat profile in $\log h_p$ (Figure 2.12(d)). These observations confirm that a Gaussian approximation is appropriate for T_A , but not for $\log h_p$.

Such graphical analysis of the negative log-likelihood is commonly referred to as “profile likelihood inspection” (see, e.g., [113]).

Accordingly, the mixed approximation strategy introduced in Section 2.5.1.5.1 is adopted: the posterior PDF $\rho_{post}(\mathbf{v})$ is approximated as the product of a Gaussian distribution for T_A and a uniform distribution for $\log h_p$.

In detail:

- the Gaussian PDF for T_A (Figure 2.13(a)) is centered at $T_A = 1341^\circ\text{C}$, with standard deviation $\sigma_{T_A,post}$ given by the square root of the (1, 1) entry of the covariance matrix Σ_{post} (see Equation (2.22)), yielding $\sigma_{T_A,post} \approx 13^\circ\text{C}$;
- the uniform posterior for $\log h_p$ is defined over the interval $(-5; -1.5)$, corresponding to the region where the likelihood remains approximately flat, as indicated by Figure 2.12(d).

The resulting posterior distributions are summarized in Table 2.8. Finally, the approximation $\bar{\sigma}_{MAP}^2$ of $\bar{\sigma}^2$ is verified to be accurate: the true value is $\bar{\sigma}^2 = 0.01$, while the estimate yields $\bar{\sigma}_{MAP}^2 = 0.00767$.

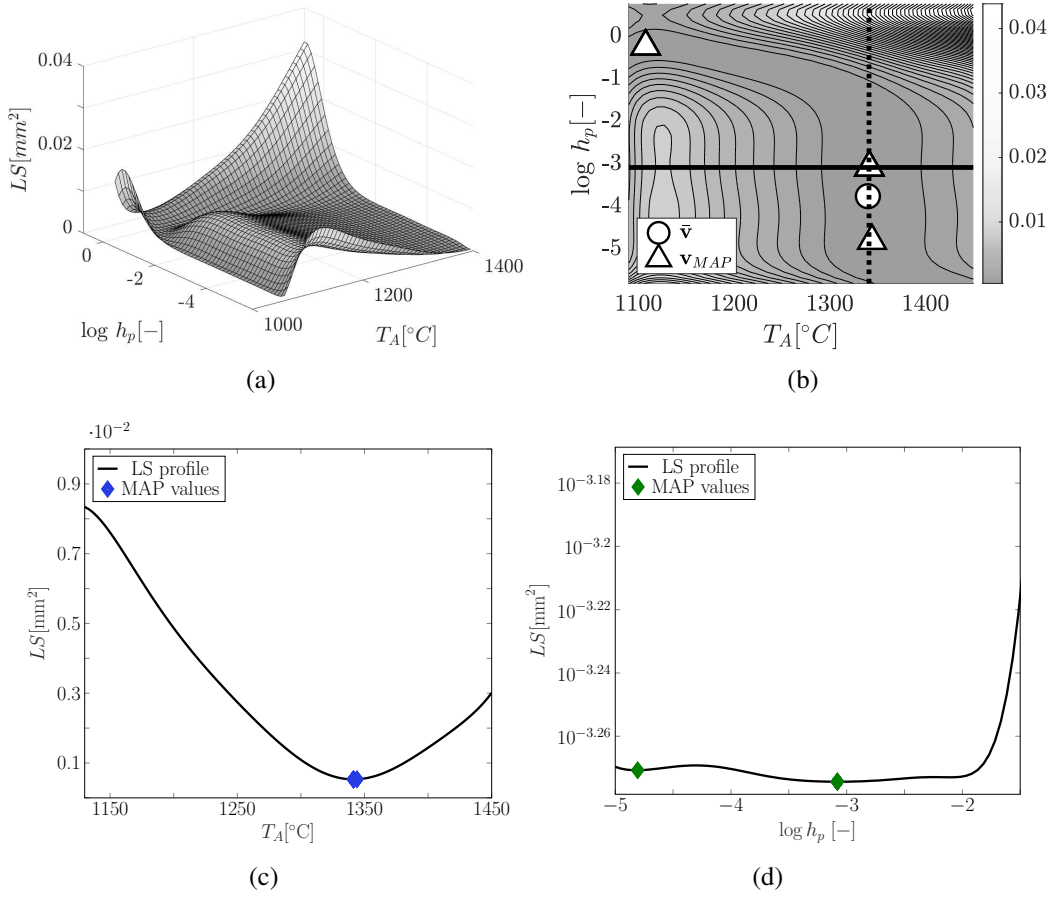


Figure 2.12: Results of the inverse UQ analysis: (a) Surface plot of the least-squares functional $LS(\mathbf{v})$; (b) Isolines of $LS(\mathbf{v})$, target value \bar{v} , the position of the cutting planes used to generate the one-dimensional plots in Figure 2.12(c) marked with a solid black line and in Figure 2.12(d) marked with a dashed black line, and the two MAP values computed by minimization of $LS(\mathbf{v})$; (c) LS profile at $\log h_p = -3$ with the two MAP values; (d) LS profile at $T_A = 1341^\circ\text{C}$ with the two MAP values.

Table 2.8: PDF of the parameters resulting from the inverse UQ analysis and used as input for the data-informed forward UQ analysis.

Random Variables	posterior PDF
T_A [°C]	Gaussian(1341; 13)
$\log h_p$ [-]	Uniform(-5, -1.5)

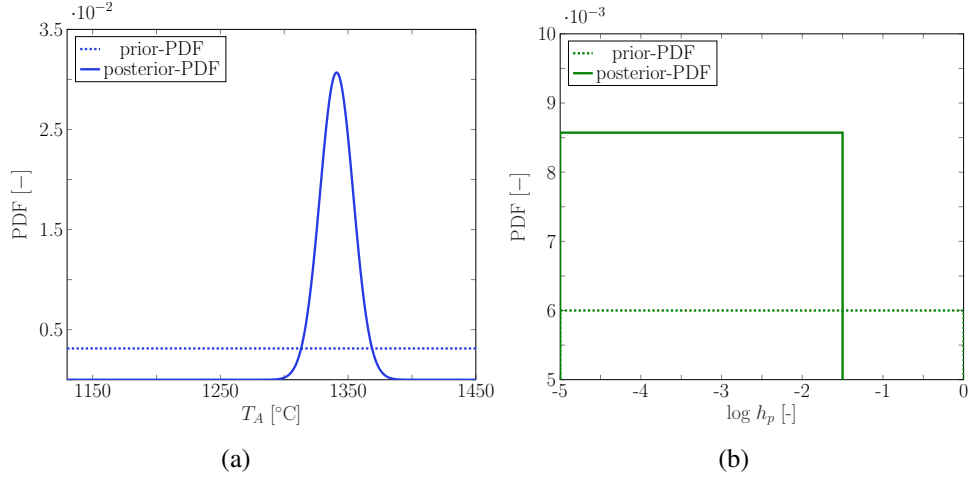


Figure 2.13: Results of the inverse UQ analysis: (a) Uniform prior-PDF and Gaussian posterior-PDF for parameter T_A before and after Bayesian inversion; (b) Uniform PDFs for the $\log h_p$ parameter before and after Bayesian inversion.

2.5.1.7.3 Data-informed forward uncertainty quantification

The final step of the uncertainty quantification workflow consists of a data-informed forward analysis. The focus is on the residual strains $\varepsilon_{xx}(\mathbf{x}_{j, str}, \mathbf{v})$ evaluated at $L = 120$ locations $\mathbf{x}_{j, str}$ in the central plane of the cantilever beam at $z = 11$ mm (see the dotted green line in Figure 2.2).

The quantities of interest are obtained by propagating the uncertainty in the parameter vector \mathbf{v} , which is sampled according to the posterior PDF ρ_{post} resulting from the inverse UQ analysis (see Table 2.8).

2.5.1.7.3.1 Surrogate modeling

Since the posterior PDF of the parameters has changed (Table 2.8), a new set of 120 sparse-grid surrogate models is constructed, following row 3 of Table 2.7. This choice reflects the need to accurately approximate the model response in the region of the parameter space that is effectively explored in the forward uncertainty propagation, which is now governed by the posterior distribution rather than the prior one.

The sparse grid underlying these surrogate models consists of 25 new collocation points in Γ_{red} and is shown in Figure 2.14(a). Figure 2.14(b) shows the sparse-grid surrogate model for the strains at $\mathbf{x} = (x, y, z) = (1.5, 2.5, 11)$ mm; a similar interpolating surface is observed for all other locations.

A convergence test similar to the one used for the sparse-grid surrogate model for the inverse UQ is also performed. Therefore, the strains are evaluated by the high-fidelity model for $M = 50$ new random values of $\mathbf{v} = (T_A, \log h_p)$ sampled according to ρ_{post} , and these residual strains are compared with their approximations obtained by the sparse-grid surrogate models with $w = 0, 1, 2, 3$, obtaining the corresponding values for the pointwise prediction error E_{PPE} (Equation (2.24)) and the quantity E_{MSE} (Equation (2.25)), which represents a normalized root

mean square error.

As expected, the trend of the errors E_{PPE} and E_{MSE} as w increases is similar for all 120 positions; therefore, the result for $\mathbf{x}_{1,str} = (1.5, 2.5, 11)$ mm is reported (see Figure 2.15). The convergence test indicates that the surrogate model with $w = 3$ (25 sparse-grid points) provides sufficient accuracy for the intended analysis.

The same conclusion is supported by Figure 2.16, which shows that, as the number of sparse-grid points increases, the residual strains predicted by the surrogate model converge to those obtained from the part-scale thermomechanical simulations.

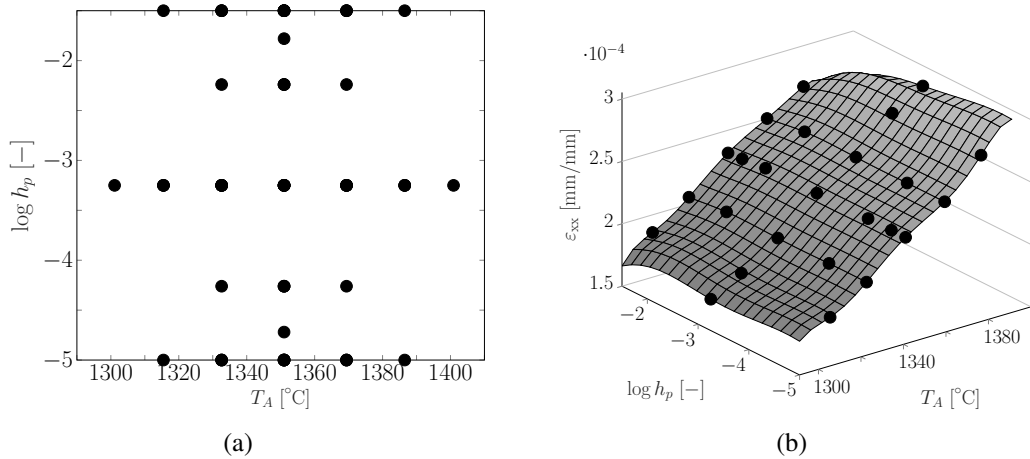


Figure 2.14: Sparse-grid surrogate model construction for the data-informed forward UQ analysis with $w = 3$ (25 sparse grid points) for the first position of the cantilever beam (i.e., $\mathbf{x}_{1,str} = (1.5, 2.5, 11)$): (a) Sparse grid; (b) Surrogate model.

2.5.1.7.3.2 Data-informed forward analysis

After validating the sparse-grid surrogate model, the final step of the data-informed forward UQ analysis is carried out, i.e., computing the data-informed PDF of ε_{xx} , as explained in Section 2.5.1.6. To assess the extent to which the inverse UQ process reduces the uncertainty in the prediction of ε_{xx} , the forward UQ procedure based on the prior-information only is also performed, i.e., a surrogate model for $\mathcal{S}_{I_{sum}} \varepsilon_{xx}$ is built according to the prior PDFs, T_A and $\log h_p$ are sampled from such prior PDF and the corresponding prior-based PDF of ε_{xx} is derived. In Figure 2.17 the following are shown: i) the most probable residual strain profiles along the x -direction x -profiles of ε_{xx} obtained by the two forward UQ analyses, i.e., the modes of the two PDFs of ε_{xx} at each of the $L = 120$ locations (dotted black line for the prior-based PDF, continuous red line for the data-informed PDF), and ii) the associated uncertainty bands, i.e., the 5%–95% quantile bands of the two PDFs (gray area for the prior-based PDF, pink area for the data-informed PDF). The figure also reports (continuous black line) the x -profile of the residual strains obtained from the part-scale thermomechanical analysis at the target value $\bar{\mathbf{v}} = (T_A; \log h_p) = (1339.8 \text{ } ^\circ\text{C}; -3.75)$: this profile is in very close agreement with the mode of




Figure 2.15: Results of the sparse-grid surrogate model for data-informed forward UQ analysis. Convergence test for the first position of the cantilever beam (i.e., $\mathbf{x}_{1, str} = (1.5, 2.5, 11)$): (a) Pointwise prediction error E_{PPE} ; (b) Root mean square error E_{MSE} .

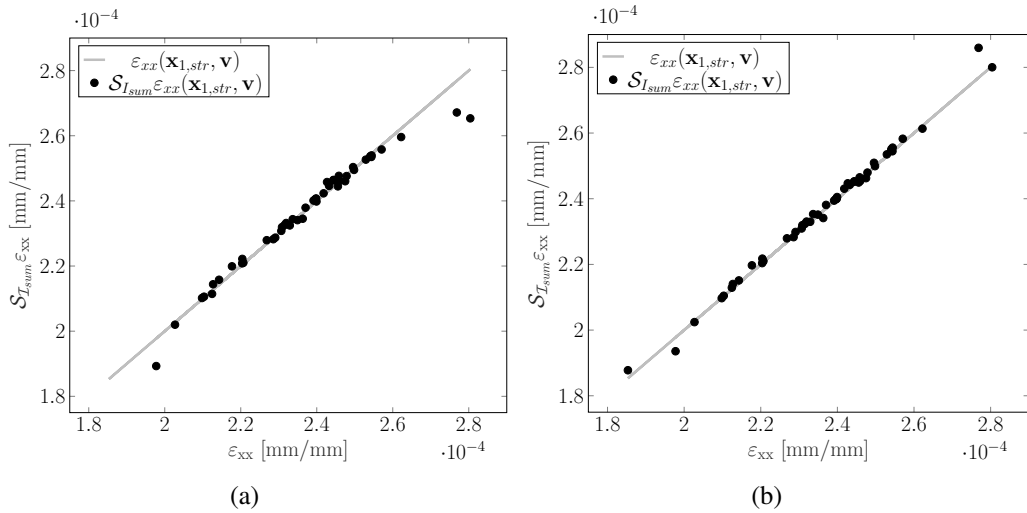


Figure 2.16: Results of the sparse-grid surrogate model for data-informed forward UQ analysis. Comparison between sparse-grid surrogate model residual strains and part-scale thermomechanical model residual strains for the first position of the cantilever beam (i.e., $\mathbf{x}_{1, str} = (1.5, 2.5, 11)$): (a) Sparse grid with level $w = 1$; (b) Sparse grid with level $w = 3$.

the data-informed PDF, which means that the most likely residual strains profile identified by such PDF closely resembles the true profile. Moreover, the fact that the prior-based quantile band is significantly larger than the data-informed one indicates that incorporating the posterior information greatly reduces the uncertainty in the prediction of the residual strains. To provide

further insight in this uncertainty reduction, 6 locations of the cantilever beam (marked by vertical dotted lines in Figure 2.17) are selected, for which the prior-based and data-informed PDFs of the residual strain are plotted in Figure 2.18. The residual strain profile for the target value $\bar{\mathbf{v}}$ is also shown in Figure 2.19, where the profile is overlapped with the geometry of the cantilever beam to provide more geometrical context to the information discussed in Figures 2.17 and 2.18.

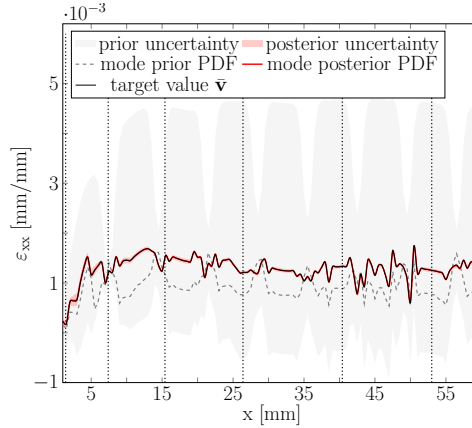


Figure 2.17: Results of data-informed forward UQ analysis: i) the mode of the prior-based (dashed black line) and data-informed posterior (continuous red line) PDFs of the residual strain ε_{xx} ; ii) the prior-based (gray area) and data-informed posterior (pink area) uncertainty bands for the residual strain ε_{xx} ; iii) the residual strain profile (continuous black line) obtained from part-scale thermomechanical analysis for the target value $\bar{\mathbf{v}}$; The vertical dotted lines represent the 6 positions at which the prior and data-informed posterior PDFs are reported in Figure 2.18.

2.5.2 Uncertainty quantification based on multi-fidelity sparse-grid surrogate modeling

The present section extends the UQ framework discussed in Section 2.5.1 for the thermomechanical response of the part-scale PBF-LB/M model. In particular, the single-fidelity surrogate modeling strategy is replaced by a multi-fidelity surrogate modeling approach based on the MISC method [81, 82].

2.5.2.1 Multi-fidelity uncertainty quantification framework

The multi-fidelity UQ framework developed in the present section follows the general structure introduced in Section 2.5.1. To avoid unnecessary repetition, all notation and assumptions concerning the uncertain parameters — including the parameter vector \mathbf{v} , the parameter domains Γ_n , the parameter space Γ , and the PDFs — are inherited directly from the single-fidelity UQ framework in Section 2.5.1. The discussion therefore focuses exclusively on the modifications introduced by the transition from a single-fidelity to a multi-fidelity setting.

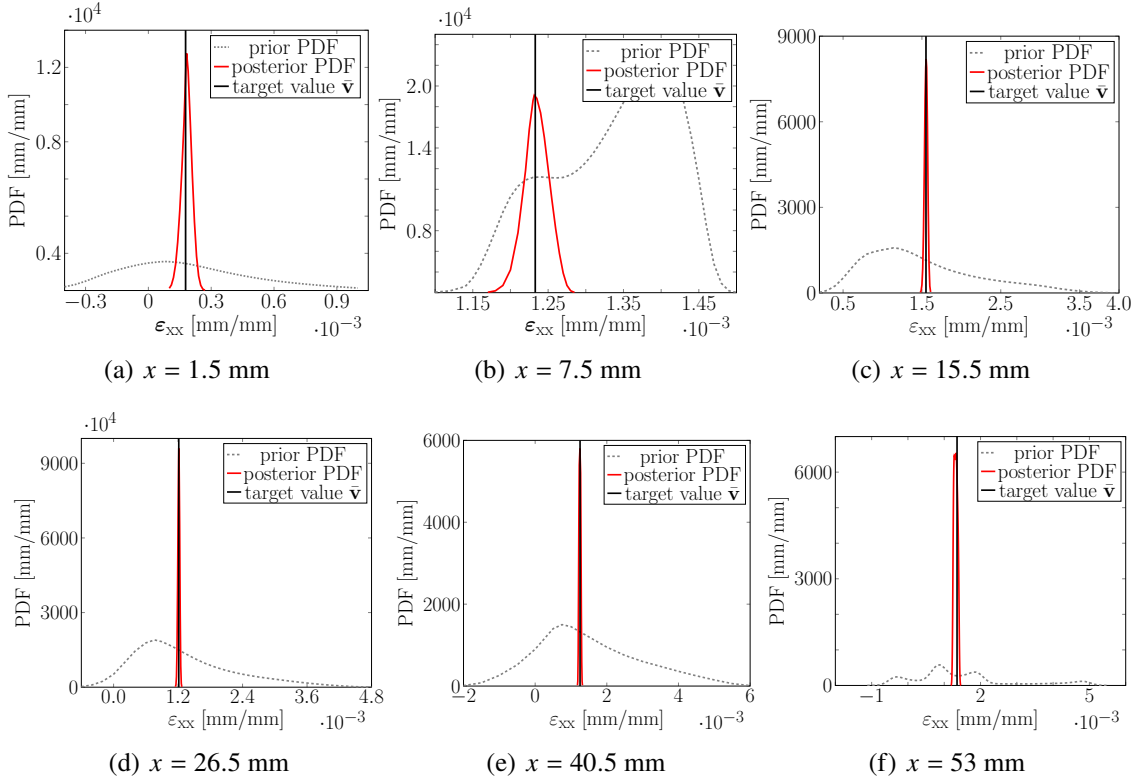


Figure 2.18: Results of data-informed forward UQ analysis. Prior-based and data-informed posterior PDFs of the residual strains at selected locations along the beam (see Figure 2.17), together with the corresponding residual strain values obtained from the part-scale thermomechanical simulation at the target value \bar{v} .

In contrast to the single-fidelity formulation, no additional GSA is performed in the multi-fidelity UQ framework. In fact, the GSA conducted in Section 2.5.1.4 identified the activation temperature, T_A , and the powder convection coefficient, h_p , as the only influential parameters, while the gas convection coefficient, h_g , exhibited a negligible impact on the QoIs, i.e., on cantilever beam displacements. Accordingly, the multi-fidelity UQ framework operates directly on the T_A – $\log h_p$ parameter space (see Table 2.6), thereby avoiding unnecessary computational cost.

A further distinction with respect to the single-fidelity UQ framework concerns the nature of the observational data employed in the Bayesian inversion. While the single-fidelity framework relied on noisy synthetic displacement data to isolate and assess the methodological aspects of the inverse problem, the multi-fidelity UQ framework directly incorporates the experimental displacement measurements provided by the NIST for the AMBench2018-01 benchmark [1, 114].

A key distinction with respect to the single-fidelity UQ framework also concerns the computational treatment of the PBF-LB/M model. High-fidelity numerical simulations, particularly those based on refined meshes required for accurate residual-strain prediction, entail a substantial com-




Figure 2.19: Residual strains obtained from the part-scale thermomechanical analysis for the target value \bar{v} .

putational expense. Since both Bayesian inverse and data-informed forward UQ analyses require repeated model evaluations across the parameter space, reliance on high-fidelity simulations alone would be computationally prohibitive.

To overcome this limitation, the present work adopts a multi-fidelity sparse-grid surrogate modeling strategy that combines a limited number of high-fidelity numerical simulations with a larger set of inexpensive low-fidelity evaluations. This approach enables Bayesian inverse and data-informed forward uncertainty quantification at a markedly reduced computational cost, while preserving the level of accuracy required for reliable thermomechanical predictions.

2.5.2.2 Multi-fidelity sparse-grid surrogate modeling

The multi-fidelity sparse-grid surrogate models are constructed by means of the MISC method, see [81, 82, 115, 116], which is a multi-fidelity surrogate modeling technique that extends to the multi-fidelity paradigm the sparse grids method, one of the well-established single-fidelity surrogate modeling methodologies [78, 80, 92, 93].

To describe how MISC works, let $f(\mathbf{v})$ denote any QoI of the PBF-LB/M model, such as the displacements or the residual strains of the cantilever beam at a certain location. The values of these QoIs can only be known approximately through the numerical scheme introduced in Section 2.4, used to solve the governing equations of the PBF-LB/M model on either of the two meshes illustrated in Figure 2.3 and detailed in Table 2.2. The notation $f_\alpha(\mathbf{v})$, with $\alpha = 1, 2$, is then introduced to refer to the value of any QoI computed on those two meshes, where $\alpha = 1$ refers to the coarse mesh and $\alpha = 2$ to the fine one. The approximations $f_\alpha(\mathbf{v})$ are referred to as *fidelties* in the following. Before continuing, it is remarked that the multi-fidelity approach described below can be extended easily to the case where more than two fidelities are available, as well as to situations in which there is more than one discretization parameter controlling the resolution of the numerical scheme (for instance, a mesh size and a time step). Furthermore, the

different fidelities may also correspond to different descriptions of the physics of the problem: for instance, distinguishing whether the physics of the printing process is simulated in a weakly coupled way (as detailed in Section 2.4) or in a fully coupled way (see, e.g., [117]). For a thorough discussion on multi-fidelity methods, see, e.g., [118, 119].

The basic idea of multi-fidelity surrogate modeling is to build an approximation of the function $\mathbf{v} \mapsto f(\mathbf{v})$ given a set of M evaluations of $f_\alpha(\mathbf{v})$ obtained for suitable choices of \mathbf{v} and α , where crucially most of the M evaluations are obtained using the low-fidelity approximation corresponding to $\alpha = 1$, thus keeping the computational cost required for the construction of the surrogate model at an acceptable level. Roughly speaking, the evaluations of $f_1(\mathbf{v})$ are needed to capture the trend of the function $\mathbf{v} \mapsto f(\mathbf{v})$, whereas the few evaluations of $f_2(\mathbf{v})$ are only needed to correct the bias in the numerical values of the surrogate model introduced by having sampled the low-fidelity approximation only. Therefore, it is evident that the success of the method largely depends on a sufficient correlation between the fidelities involved. The presence of such correlation is not always guaranteed and may be less common than sometimes believed; one situation in which correlation is present is when the different fidelities consist of different computational meshes, all of which are sufficiently accurate to capture the essence of the physics (as in the present work). In detail, the MISC multi-fidelity surrogate model of $f(\mathbf{v})$ is computed as a linear combination of several single-fidelity surrogate models of $\mathbf{v} \mapsto f_\alpha(\mathbf{v})$, each obtained by evaluating f_α over a different set of values of \mathbf{v} , according to the following equation:

$$f(\mathbf{v}) \approx \mathcal{S}_I f(\mathbf{v}) = \sum_{[\alpha, \beta] \in \mathcal{I}} c_{\alpha, \beta} \mathcal{U}_{\alpha, m(\beta)}(\mathbf{v}). \quad (2.26)$$

In the equation above:

1. $m(\cdot)$ is a non-decreasing function; in the present work the following definition is considered, although other choices are possible (see, e.g., [80]):

$$m : \mathbb{N}_+ \rightarrow \mathbb{N}_+ \text{ such that } m(i) = 2i - 1.$$

2. $\beta = (\beta_1, \dots, \beta_N) \in \mathbb{N}_+^N$ is a multi-index whose components control the sampling of the parameter space Γ . More specifically, a Cartesian grid over Γ , denoted $\mathcal{T}_m(\beta)$, is considered, composed of $m(\beta_1) \times m(\beta_2) \times \dots \times m(\beta_N)$ points. The sets of univariate points used to generate $\mathcal{T}_m(\beta)$ are denoted by $\mathcal{T}_{n, m(\beta_n)}$. For efficiency reasons, the sets $\mathcal{T}_{n, m(\beta_n)}$ should be chosen according to the PDF ρ_n , and should satisfy $\mathcal{T}_{n, m(r)} \subset \mathcal{T}_{n, m(s)}$ if $r < s$, i.e., the collocation points must be nested. A good choice for ρ_n uniform, as in the present problem, is represented by the so-called symmetric Leja points as in Appendix B. The Cartesian grid $\mathcal{T}_m(\beta)$ is then the set of points $\mathcal{T}_m(\beta) = \left\{ \mathbf{v} \in \Gamma : v_n = v_{n, m(\beta)}^{(j_n)}, j_n \leq m(\beta_n) \right\}$. It is already noted that after the Bayesian inverse analysis the PDFs ρ_n of the uncertain parameters will no longer be uniform, so these points will need to be changed and the multi-fidelity sparse-grid surrogate model will need to be re-built.
3. $\mathcal{U}_{\alpha, m(\beta)}(\mathbf{v})$ is an N -variate Lagrangian interpolant of $f_\alpha(\mathbf{v})$ built over the Cartesian grid

$\mathcal{T}_{m(\boldsymbol{\beta})}$, i.e.,

$$f_\alpha(\mathbf{v}) \approx \mathcal{U}_{\alpha, m(\boldsymbol{\beta})}(\mathbf{v}) := \sum_{\mathbf{j} \leq m(\boldsymbol{\beta})} f_\alpha \left(v_{m(\boldsymbol{\beta})}^{(\mathbf{j})} \right) \mathcal{L}_{m(\boldsymbol{\beta})}^{(\mathbf{j})}(\mathbf{v}),$$

where $\mathbf{j} \leq m(\boldsymbol{\beta})$ means $j_n \leq m(\beta_n)$ for $n = 1, \dots, N$, and $\mathcal{L}_{m(\boldsymbol{\beta})}^{(\mathbf{j})}$ are the N -variate Lagrange polynomials associated with the points in $\mathcal{T}_{m(\boldsymbol{\beta})}$.¹

4. $\mathcal{I} \subset \mathbb{N}_+^{1+N}$ is a multi-index set whose elements are extended multi-indices $[\alpha, \boldsymbol{\beta}]$. It gathers all the N -variate Lagrangian interpolants that compose the MISC approximation. It is required that \mathcal{I} is downward closed, i.e.,

$$[\gamma, \boldsymbol{\delta}] \leq [\alpha, \boldsymbol{\beta}] \text{ and } [\alpha, \boldsymbol{\beta}] \in \mathcal{I} \Rightarrow [\gamma, \boldsymbol{\delta}] \in \mathcal{I}.$$

The performance of the multi-fidelity sparse-grid surrogate model based on MISC (i.e., its accuracy as the number of collocation points increases) crucially depends on the choice of the multi-index set \mathcal{I} . In the present work, an adaptive a-posteriori strategy is considered to simultaneously enlarge the multi-index set \mathcal{I} and refine the corresponding surrogate model $\mathcal{S}_{\mathcal{I}} f(\mathbf{v})$, see [115, 116] for details. Roughly speaking, at each iteration of such adaptive algorithm:

- (a) a few candidate extended multi-indices $[\alpha, \boldsymbol{\beta}]$ are temporarily added to \mathcal{I} ;
 - (b) the corresponding new surrogate model $\mathcal{S}_{\mathcal{I}} f(\mathbf{v})$ is computed, which requires solving the PBF-LB/M model equation on the mesh specified by α for the parameter values dictated by $\boldsymbol{\beta}$;
 - (c) a heuristic selection criterion selects the extended multi-index $[\alpha, \boldsymbol{\beta}]$ that is most effective in improving the surrogate model [115, 116];
 - (d) the list of candidates is updated accordingly, and the next iteration begins.
5. $c_{\alpha, \boldsymbol{\beta}}$ are the so-called combination technique coefficients defined as follows:

$$c_{\alpha, \boldsymbol{\beta}} := \sum_{\substack{i \in \{0,1\}, \mathbf{j} \in \{0,1\}^N: \\ [\alpha+i, \boldsymbol{\beta}+\mathbf{j}] \in \mathcal{I}}} (-1)^{\|\mathbf{j}\|_1}.$$

2.5.2.3 Results of the uncertainty quantification workflow

In the present section, the numerical results obtained while performing the multi-fidelity UQ workflow are shown and discussed. In detail, a procedure to compute the data-informed PDF of the uncertain parameters by Bayesian inverse techniques is first presented, and the data-informed forward UQ analysis for the residual strains is then illustrated.

¹i.e., $\mathcal{L}_{m(\boldsymbol{\beta})}^{(\mathbf{j})}(\mathbf{v}) = \prod_{n=1}^N \ell_{n, m(\beta_n)}^{(j_n)}(v_n)$ with $\ell_{n, m(\beta_n)}^{(j_n)}(v_n) = \prod_{k=1, k \neq j_n}^{m(\beta_n)} \frac{v_n - v_{n, m(\beta_n)}^{(k)}}{v_{n, m(\beta_n)}^{(k)} - v_{n, m(\beta_n)}^{(j_n)}}$.

2.5.2.3.1 Inverse uncertainty quantification analysis

2.5.2.3.1.1 Uncertain parameters and available data

As discussed in Section 2.5.2.1 the analysis considers the activation temperature, T_A , and the base-10 logarithm of the powder convection coefficient, $\log h_p$.

The experimental measurements data considered to perform the Bayesian inverse UQ analysis are the displacements at the first $K = 5$ ridges of the cantilever beam provided by NIST [1], denoted as $\mathbf{x}_{k,meas}$, $k = 1, \dots, K$ (see Figure 2.20). Only the first 5 ridges of the cantilever beam are considered, as further tests with a larger number of data did not change the essence of the results shown below. It is assumed that such experimental data correspond to the solution of the PBF-LB/M model for an unknown value of the uncertain parameters \mathbf{v}_{true} , corrupted by an experimental error distributed as a Gaussian random variable with zero mean and standard deviation σ_{meas} ; furthermore, the experimental errors at the K locations are statistically independent. In mathematical terms:

$$u_{k,exp} = u(\mathbf{x}_{k,meas}, \mathbf{v}_{true}) + \epsilon_k, \quad \epsilon_k \sim \mathcal{N}(0, \sigma_{meas}^2), \quad k = 1, \dots, K. \quad (2.27)$$

The values of $u_{k,exp}$, $k = 1, \dots, K$ are reported in Figure 2.21, together with the displacements at the 6 additional ridges that are not considered, as already motivated.




Figure 2.20: PBF-LB/M numerical model of the cantilever beam 75 mm long, 12 mm high and 5 mm wide with a build platform measuring 85 mm long, 12 mm high and 20 mm wide. Points marked in purple are the locations $\mathbf{x}_{k,meas}$, $k = 1, \dots, 5$ used during the Bayesian inverse UQ analysis; the (blue) dotted line marks the locations where residual strains are predicted using data-informed forward UQ analysis.

2.5.2.3.1.2 Surrogate modeling

The MISC sparse-grid surrogate model considered in the present analysis uses 22 evaluations of the PBF-LB/M numerical model, corresponding to the points in the parameter space represented in Figure 2.22(a). In detail, 17 evaluations of the PBF-LB/M numerical model are performed




Figure 2.21: Displacements at the 11 ridges at the top of the cantilever beam provided by NIST [1].

at the low fidelity, $\alpha = 1$, and 5 at the high fidelity, $\alpha = 2$. Note that the 5 points for which a high-fidelity numerical simulation is requested are a subset of the 17 points for the low-fidelity simulation, due to the nestedness of the symmetric Leja points.

The surrogate models obtained for the K locations $\mathbf{x}_{k,meas}$ show similar behaviour; Figure 2.22(b) reports the one obtained for the first ridge of the cantilever beam, with coordinates $x = 0.5$ mm, $y = 2.5$ mm, $z = 12.5$ mm.

To further clarify this choice, it is important to recall that low fidelity is associated with the coarse mesh of the cantilever beam model, whereas high fidelity is associated with the fine mesh, as discussed in Section 2.4.1. When switching from low fidelity to high fidelity, the mesh size is doubled, i.e., the number of elements in the vertical direction of the cantilever beam model (Figure 2.3) is increased, as shown in Table 2.2. This choice is motivated by the aim of accurately capturing the 5 ridges at the top of the cantilever beam in both models, as illustrated in Figure 2.3. The adoption of two levels of fidelity in the cantilever beam model is related to the fact that, although the low-fidelity model is able to represent the PBF-LB/M process, it cannot capture all the details of the real geometry, in particular the openings in the cantilever beam (see Figure 2.1). These details are captured by the high-fidelity numerical model. Therefore, using the MISC methodology, the surrogate model is constructed by exploiting the low-fidelity model to replicate the PBF-LB/M process and then correcting it with the high-fidelity model, which is crucial for more accurate results.

The adoption of the multi-fidelity sparse-grid surrogate model based on MISC with 17 evaluations on the low-fidelity model and 5 evaluations on the high-fidelity model is driven by a performance analysis of the surrogate model. The adaptive MISC algorithm detailed in Section 2.5.2.2 is executed until a surrogate model providing acceptable pointwise prediction error and mean square error values over a specified set of validation points is obtained. These validation points are deliberately excluded from the surrogate construction and are used solely for validation purposes, in order to provide an unbiased estimate of the prediction error. Including them in




Figure 2.22: MISC surrogate model construction based on 17 evaluations (points marked in blue) of the low-fidelity PBF-LB/M numerical model ($\alpha = 1$) and 5 evaluations (points marked in yellow) of the high fidelity ($\alpha = 2$) for the first node of the cantilever beam (i.e., $\mathbf{x}_{1,meas} = (0.5, 2.5, 12.5)$): (a) sparse grid; (b) surrogate model.

the surrogate construction would compromise the independence of the validation procedure and could lead to an overestimation of the surrogate accuracy. To evaluate the effectiveness of the obtained surrogate model, the displacement results obtained through numerical simulations with *ANSYS2021-R2* software on the high-fidelity model, U_Z , are compared with the results interpolated from the multi-fidelity sparse-grid model, $\mathcal{S}_I U_Z$, for 7 pairs of uncertain parameters (T_A, h_p) values (validation points) represented by the burgundy points in Figure 2.9. In particular, Figure 2.9 compares the displacement results using: the MISC model based on 6 evaluations (5 on the low-fidelity model and 1 on the high-fidelity model), Figure 2.23(a); the MISC model based on 10 evaluations (7 on the low-fidelity model and 3 on the high-fidelity model), Figure 2.23(b); and the MISC model based on 22 evaluations (17 on the low-fidelity model and 5 on the high-fidelity model), Figure 2.23(c). The analysis of the results reveals a remarkable degree of accuracy in the selected multi-fidelity surrogate model (Figure 2.23(c)). Indeed, contrary to what is shown in Figures 2.23(a) and 2.23(b), Figure 2.23(c) shows that all points interpolated on the multi-fidelity surrogate model exhibit significant alignment with the results obtained from *ANSYS2021-R2* (represented by the blue line), indicating a good predictive ability of the surrogate model. This highlights the importance of an adequate number of evaluations to ensure the robustness and reliability of the multi-fidelity sparse-grid surrogate model based on MISC.

This level of consistency between the results obtained from the direct simulations and those interpolated from the multi-fidelity surrogate model suggests that the multi-fidelity surrogate model can effectively capture the complex relationships between the input parameters and the responses of the thermomechanical cantilever beam model. Interestingly, this accuracy is achieved despite the multi-fidelity surrogate model being trained using a relatively small number

of numerical evaluations, thus demonstrating the effectiveness of the MISC method in optimizing the balance between accuracy and computational cost.

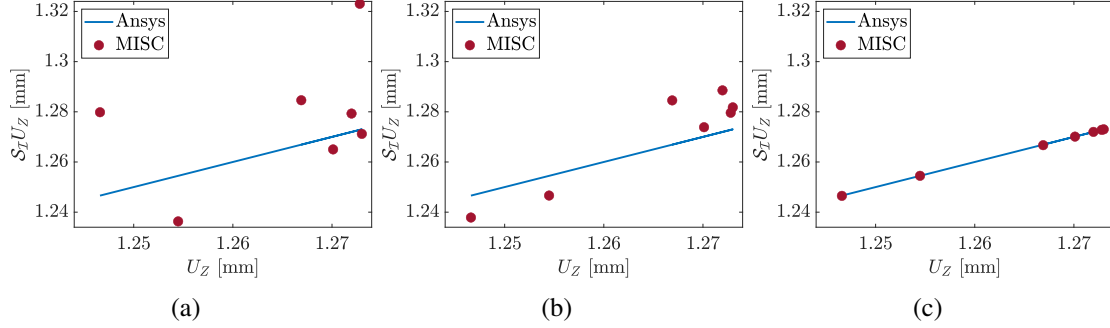


Figure 2.23: Results of the MISC surrogate model for inverse UQ analysis. Comparison between high-fidelity displacement results with ANSYS2021-R2, U_Z , and the displacement results for three different MISC surrogate models ($\mathcal{S}_I U_Z$): a) MISC model with 6 (5+1) evaluations; b) MISC model with 10 (7+3) evaluations; c) MISC model with 22 (17+5) evaluations. The blue line represents the pairs $(U_{Z,ansys}, U_{Z,ansys})$ and the burgundy points represent the pairs $(U_{Z,ansys}, U_{Z,MISC})$, evaluated for each of the 7 validation points.

2.5.2.3.1.3 Bayesian inverse analysis

As discussed in Section 2.5.1.5 the goal of the Bayesian inverse UQ analysis is to find the posterior PDF, $\rho_{post}(\mathbf{v})$, of $\log h_p$ and T_A [61, 62, 120, 121] which is approximated locally by a Gaussian distribution via a Laplace approximation. The approach requires setting the center of $\rho_{post}(\mathbf{v})$ at the MAP point of the true posterior PDF, \mathbf{v}_{MAP} , i.e.:

$$\mathbf{v}_{MAP} := \arg \max_{\mathbf{v} \in \mathbb{R}^N} \rho_{post}(\mathbf{v}), \quad (2.28)$$

where:

$$\rho_{post}(\mathbf{v}) = \mathcal{L}(\mathbf{v} | u_{1,exp}, u_{2,exp}, \dots) \rho_{prior}(\mathbf{v}) \frac{1}{C}.$$

More specifically, under the assumption made in Section 2.5.4.1.1 that the experimental errors ϵ_k are statistically independent random variables with PDF ρ_{ϵ_k} , cf. Equation (2.27), the likelihood function can be written as:

$$\mathcal{L}(\mathbf{v} | u_{1,exp}, u_{2,exp}, \dots) = \prod_{k=1}^5 \rho_{\epsilon_k}(u_{k,exp} - u(\mathbf{x}_{k,meas}, \mathbf{v})).$$

Given the expression above for the likelihood, and under the further assumptions that the prior PDF of the uncertain parameters is independently uniform (i.e., $\rho_{prior}(\mathbf{v})$ is a constant over Γ and zero elsewhere, cf. Section 2.5.2.1), and that the experimental errors are Gaussian random

variables, cf. Equation (2.27), it can be easily shown [61, 62, 120, 121] that the maximisation problem in Equation (2.28) is in practice equivalent to:

$$\begin{aligned} \mathbf{v}_{MAP} &= \arg \min_{\mathbf{v} \in \Gamma} [-\log \mathcal{L}(\mathbf{v} | u_{1,exp}, u_{2,exp}, \dots)] = \arg \min_{\mathbf{v} \in \Gamma} \sum_{k=1}^K (u_{k,exp} - u(\mathbf{x}_{k,meas}, \mathbf{v}))^2 \quad (2.29) \\ &\approx \arg \min_{\mathbf{v} \in \Gamma} \sum_{k=1}^K (u_{k,exp} - \mathcal{S}_I u(\mathbf{x}_{k,meas}, \mathbf{v}))^2, \end{aligned}$$

where crucially in the last step the high-fidelity PBF-LB/M model evaluations $u(\mathbf{x}_{k,meas}, \mathbf{v})$ have been replaced with their multi-fidelity surrogate model approximations to speed up the computations.

Also, it is noted that the minimisation in Equation (2.29) does not depend on the standard deviation of the experimental errors, σ_{meas} . If needed, a sample estimate of this quantity can be obtained as:

$$\sigma_{meas}^2 \approx \frac{1}{K} \sum_{k=1}^K (u_{k,exp} - \mathcal{S}_I u(\mathbf{x}_{k,meas}, \mathbf{v}_{MAP}))^2.$$

Once \mathbf{v}_{MAP} is obtained, the next step is computing the covariance matrix Σ_{post} of the Gaussian approximation of $\rho_{post}(\mathbf{v})$. Under the same assumptions that lead to Equation (2.29), it can be shown that

$$\Sigma_{post} \approx \sigma_{meas}^2 (\mathcal{J}_u^T \mathcal{J}_u)^{-1},$$

where \mathcal{J}_u is the Jacobian of the PBF-LB/M model prediction $u(\mathbf{x}_{k,meas}, \mathbf{v})$ evaluated at the MAP point, i.e.,

$$[\mathcal{J}_u]_{k,j} = \frac{\partial}{\partial v_j} u(\mathbf{x}_{k,meas}, \mathbf{v}_{MAP}).$$

Following the steps just described, the resulting estimate is $\mathbf{v}_{MAP} = (1386^\circ\text{C}; -3)$ and Σ_{post} is essentially diagonal, which means that T_A and $\log h_p$ are essentially statistically independent, i.e., two independent Gaussian variables; the diagonal entries of Σ_{post} provide their standard deviations, $\sigma_{T_A,post} = 95^\circ\text{C}$ and $\sigma_{\log h_p,post} = 0.92$.

It is important to note that for the calibration of the parameter \mathbf{v} (Equation (2.29)) the Nelder–Mead gradient-free optimization algorithm, accessible in MATLAB through the `fminsearch` command, has been used. To increase the robustness of the results, several optimizations with 20 different initial conditions for the optimization algorithm have been performed, requiring overall 10000 surrogate model evaluations.

In Figure 2.24 the prior and posterior PDFs of the two parameters are compared.

It is important to remark that the Gaussian approximation of the posterior PDF obtained via the Laplace method is a local approximation around the MAP point and does not enforce the bounded support of the prior distribution. As a consequence, the resulting Gaussian posterior may assign non-zero probability mass outside the admissible parameter domain Γ .

This behavior is visible in Figure 2.24, where a portion of the approximate posterior distributions extends beyond the prior bounds. This does not imply a violation of Bayes' theorem,




Figure 2.24: Results of the Bayesian inverse UQ analysis: (a) uniform prior PDF and Gaussian posterior PDF for parameter $\log h_p$ before and after Bayesian inverse UQ analysis; (b) uniform prior PDF and Gaussian posterior PDF for parameter T_A before and after Bayesian inverse UQ analysis.

Table 2.9: Statistical information for the two uncertain parameters before and after the Bayesian inverse. The reported intervals correspond to approximate 99.7% credible intervals defined as $\mu \pm 3\sigma$.

Prior	mean	st. dev.	CoV	interval
T_A	1290°C	92.37°C	0.072	[1130; 1450]°C
$\log h_p$	-2.5	1.44	0.577	[-5; 0]
Posterior	mean	st. dev.	CoV	interval
T_A	1386°C	95°C	0.068	[1100; 1671]°C
$\log h_p$	-3	0.92	0.306	[-5.76; -0.24]

but rather reflects the fact that the Gaussian approximation does not preserve the compact support of the exact posterior distribution.

The Gaussian approximation should therefore be interpreted as a local characterization of the posterior distribution in the vicinity of the MAP point, suitable for uncertainty propagation, rather than as a fully consistent global representation of the posterior PDF.

A more accurate representation of the posterior distribution, preserving the support of the prior, could be obtained by sampling-based methods such as Markov Chain Monte Carlo (MCMC), at the price of a significantly higher computational cost.

Note, in particular, that despite this limitation, a reduction of uncertainty is still observed when comparing prior and posterior statistics. This can be verified by comparing the standard deviations and the coefficients of variation (CoV) of the parameters reported in Table 2.9, which show a significant reduction of uncertainty on $\log h_p$ and a slight reduction for T_A .

2.5.2.3.2 Data-informed forward uncertainty quantification analysis

The aim of the data-informed forward UQ analysis is to predict the profile of the residual strains of the cantilever beam and their associated uncertainties. Specifically, the quantities of interest are the residual strains of $J = 120$ nodes located in the middle plane of the cantilever beam along the x -direction with a sampling step of 0.5 mm and for $y = 2.5$ mm and $z = 11$ mm, $\varepsilon_{xx}(\mathbf{x}_{j, str}, \mathbf{v})$ with $j = 1, \dots, J$ (Figure 2.20).

To do this, $S = 10000$ samples of \mathbf{v} are generated by sampling from ρ_{post} and for each one $\varepsilon_{xx}(\mathbf{x}_{j, str}, \mathbf{v})$ is evaluated, which again requires introducing suitable MISC surrogate models. The PDF of the residual strains at each cantilever beam node is then estimated by applying a kernel density estimation method [109, 110] to the S samples. Finally, the residual strains of the cantilever beam are predicted by evaluating the mode of these PDFs. The results obtained in this way are discussed below, and the reader is referred to Section 2.5.1 [122] for more details on the methodology.

2.5.2.3.2.1 Surrogate modeling

To now compute the residual strains of the cantilever beam, a new set of surrogate models needs to be built, again by the MISC method. Note that the PBF-LB/M simulations that were run at the collocation points in Figure 2.22(a) for building the surrogate models for the Bayesian inverse are no longer usable, and a new set of simulations is required. This is for two reasons: first, those collocation points were selected based on the assumption that uncertain parameters were uniform random variables, whereas Gaussian random variables are now used, which calls for a different set of collocation points to maintain good approximation properties (cf. Item 2 of the enumerated list in Section 2.5.2.2). Second, even if a suboptimal surrogate model were acceptable, the surrogate model used for the Bayesian inverse is by construction only valid in Γ , whereas the support of ρ_{post} significantly overflows such domain.

The new set of collocation points is shown in Figure 2.25 and uses the weighted Gaussian Leja points [123–126], which are suitable for Gaussian random variables. Initially, only 6 evaluations of the PBF-LB/M numerical model are used: 5 evaluations on the low-fidelity model corresponding to $\alpha = 1$, and one evaluation on the high-fidelity model corresponding to $\alpha = 2$: more details on the number of collocation points are provided later on. Examples of the surrogate models constructed for the J cantilever beam nodes are shown in Figure 2.26.




Figure 2.25: Collocation points used for the construction of the MISC multi-fidelity surrogate models based on 5 evaluations (points marked in blue) of the PBF-LB/M numerical model at low fidelity, $\alpha = 1$, and 1 evaluation (marked in yellow) of the PBF-LB/M numerical model at high fidelity, $\alpha = 2$ for the residual strains cantilever beam model.

2.5.2.3.2.2 Data-informed forward analysis

It is now possible to approximate the residual strains at the S random samples of T_A and $\log h_p$ and compute the corresponding PDFs. For the sake of comparison, the PDFs of the residual strains that can be obtained according to ρ_{prior} are also reported, i.e., upon building a set of surrogate models for the residual strains based on ρ_{prior} (note that the surrogate models discussed in Section 6.1.2, see, e.g., Figure 6, were for the displacement and not for the residual strains), and evaluating them on a ρ_{prior} -based sample of values for T_A and $\log h_p$.

In Figure 2.27 the results obtained are presented, which include: (i) the most likely profiles of ε_{xx} , according to ρ_{prior} and ρ_{post} respectively, i.e., the modes of the two PDFs of ε_{xx} at each of the $J = 120$ positions (the dashed grey line represents the mode of the prior PDFs, while the solid red line represents the mode of the posterior PDFs), and (ii) the associated uncertainty bands, i.e., the 5%–95% quantile bands of the two PDFs (grey area for the mode PDF based on the prior PDFs and pink area for the mode posterior-based PDFs). Figure 2.27 also shows (iii) the residual strains profile provided by NIST (continuous blue line), while the vertical dotted lines represent the 6 positions at which the prior and data-informed PDFs are reported in Figure 2.28. The residual strains profile provided by NIST overlaps the mode of the posterior-based PDFs, indicating that the most likely residual strains profile identified by these PDFs is a faithful representation of the residual strains experimental profile.

Figure 2.27 also shows that the quantile bands based on the posterior PDFs are less wide than those based on the prior PDFs, indicating that overall the Bayesian inverse procedure was able to significantly reduce the uncertainty in the predictions of the residual strains.

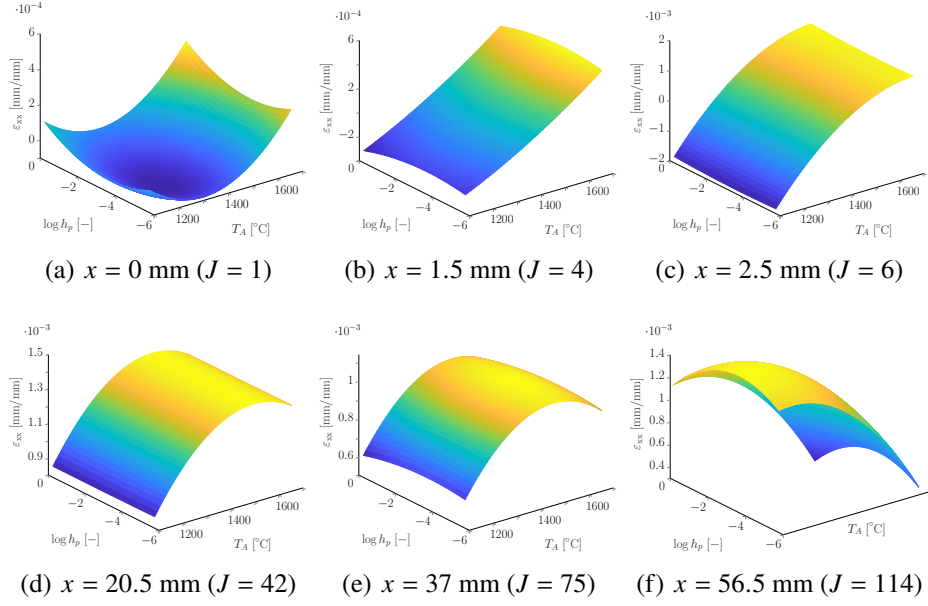


Figure 2.26: Results of the data-informed forward UQ analysis. MISC multi-fidelity surrogate models for some of the nodes of the cantilever beam, obtained from the collocation points shown in Figure 2.25.

To give a more quantitative appraisal of this statement, the quantile band width for the residual strains at each of the J positions is computed using both prior and posterior PDFs. The differences between the two quantile bands are then averaged, with normalisation relative to the prior quantile band. More precisely, the uncertainty reduction in residual strain predictions is computed as follows:

$$100 \times \frac{1}{J} \sum_{j=1}^J \frac{\left(Q_{prior,95\%}(\varepsilon_{xx,j}) - Q_{prior,5\%}(\varepsilon_{xx,j}) \right) - \left(Q_{post,95\%}(\varepsilon_{xx,j}) - Q_{post,5\%}(\varepsilon_{xx,j}) \right)}{Q_{prior,95\%}(\varepsilon_{xx,j}) - Q_{prior,5\%}(\varepsilon_{xx,j})},$$

where $Q_{prior,5\%}$ and $Q_{prior,95\%}$ represent the 5% and 95% quantile bands relative to the prior PDFs of ε_{xx} for $j = 1, \dots, J$ locations, respectively. Analogously, $Q_{post,5\%}$ and $Q_{post,95\%}$ are the 5% and 95% quantile bands relative to the posterior PDFs of ε_{xx} for $j = 1, \dots, J$ locations, respectively. The results indicate a 33% reduction from prior to posterior which is a rather significant improvement. Of course, these results are somehow biased by the fact that relatively large prior intervals for T_A and $\log h_p$ were willingly taken (cf. Section 2.5.2.3.1.1), i.e., taking narrower prior intervals might have led to a smaller reduction in the uncertainty. On the other hand, the narrower the prior intervals, the higher the risk of missing \mathbf{v}_{MAP} , i.e., the risk that $\mathbf{v}_{MAP} \notin \Gamma$ and thus the Gaussian approximation of the posterior PDF would be centered at the wrong place. These results thus overall show that one can safely start from relatively wide prior intervals, as the methodology is quite robust in this respect.

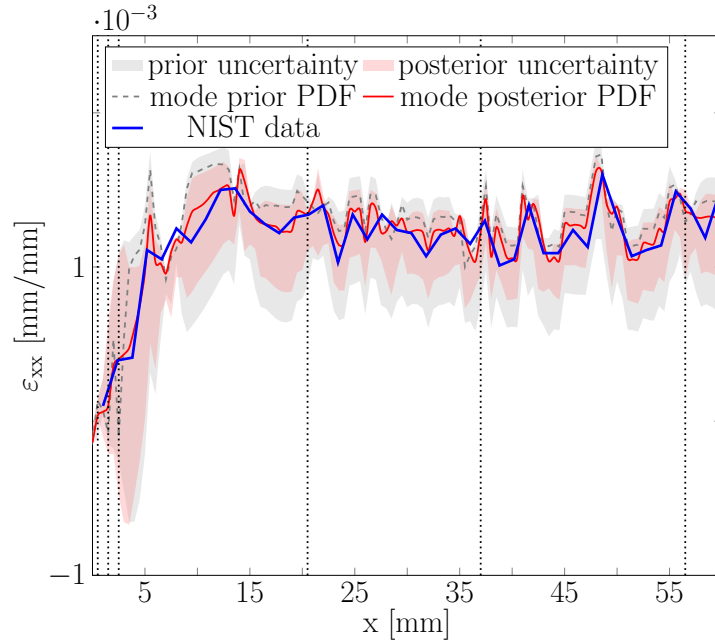


Figure 2.27: Results of the data-informed forward UQ analysis. The figure includes: (i) the most likely profiles of ε_{xx} obtained from forward UQ analyses based on ρ_{prior} (the dashed grey line represents the mode of the prior-based PDFs) and ρ_{post} (the solid red line represents the mode of the posterior-based PDFs), (ii) the associated uncertainty bands, i.e., the 5%–95% quantile bands of the two PDFs (grey area for the mode prior-based PDF and pink area for the mode posterior-based PDF) and (iii) the residual strains profile provided by NIST (continuous blue line). The vertical dotted lines represent the 6 positions at which the prior and data-informed posterior PDFs are reported in Figure 2.28.

To further check the validity of the results presented in Figure 2.27, the forward UQ analysis was repeated upon refining twice the multi-fidelity surrogate model and comparing the results thus obtained. Specifically, two new surrogate models were constructed by using 7 and 13 evaluations of the low-fidelity model (i.e., incrementally adding first 2 collocation points to those in Figure 2.25 and then 6 more, given the nestedness of the weighted Leja points — no high-fidelity evaluations are added), and the results reported in Figure 2.29 were obtained. Such results show that, despite the increase in the number of evaluations of the PBF-LB/M numerical model, there is no significant change, neither in the quantile bands nor in the modes of the PDFs of the residual strains uncertainty.

This confirms the validity of choosing the model based on 5 low-fidelity model evaluations and 1 high-fidelity model evaluation, as it is not only able to faithfully represent the experimental residual strains, but is also very cost effective from a computational point of view.

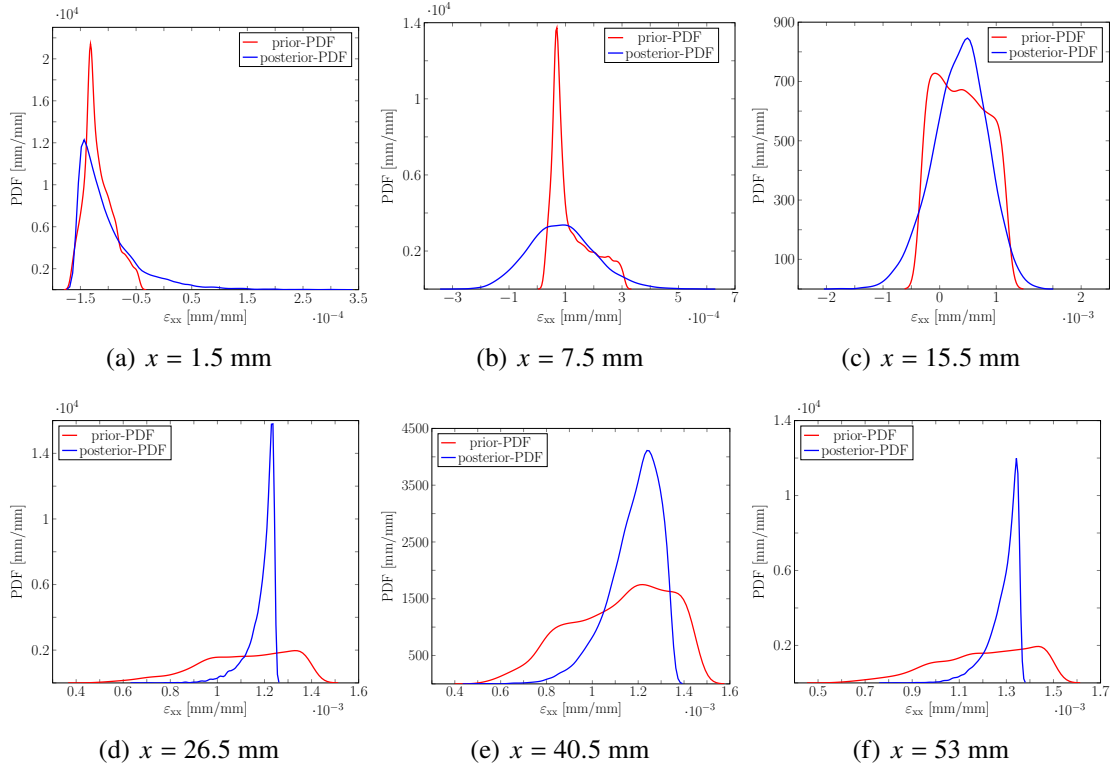


Figure 2.28: Results of data-informed forward UQ analysis. Prior-based and data-informed PDFs of the residual strains of the cantilever beam for the x locations marked in Figure 2.20.

2.5.2.3.2.3 Multi-fidelity versus high-fidelity surrogate models results

To conclude the discussion, the results obtained with the multi-fidelity surrogate model (5 low-fidelity + 1 high-fidelity) are compared with those obtained by employing a surrogate model built using high-fidelity simulations only, running 5 high-fidelity simulations at the same collocation points where the multi-fidelity surrogate model requests low-fidelity simulations.

The results in Figure 2.30 show that the predicted modes obtained with the two approaches are qualitatively similar in the region of $x > 25$ mm, while differences can be observed in the associated uncertainty bands. In particular, the multi-fidelity uncertainty band appears below the reference data, whereas the high-fidelity band is shifted upward. For $x < 25$ mm, the multi-fidelity model exhibits a wider uncertainty band compared to the high-fidelity model.

It is important to remark that this comparison is based on a high-fidelity surrogate model constructed from a very limited number of simulations. Therefore, the observed differences should be interpreted with caution and primarily provide a qualitative indication of the behavior of the two approaches, rather than a definitive assessment of their relative accuracy.

On the other hand, the computational time required to run 5 high-fidelity simulations is 36 times larger than the time of 5 low-fidelity analyses (see Table 2.2). This highlights the significant

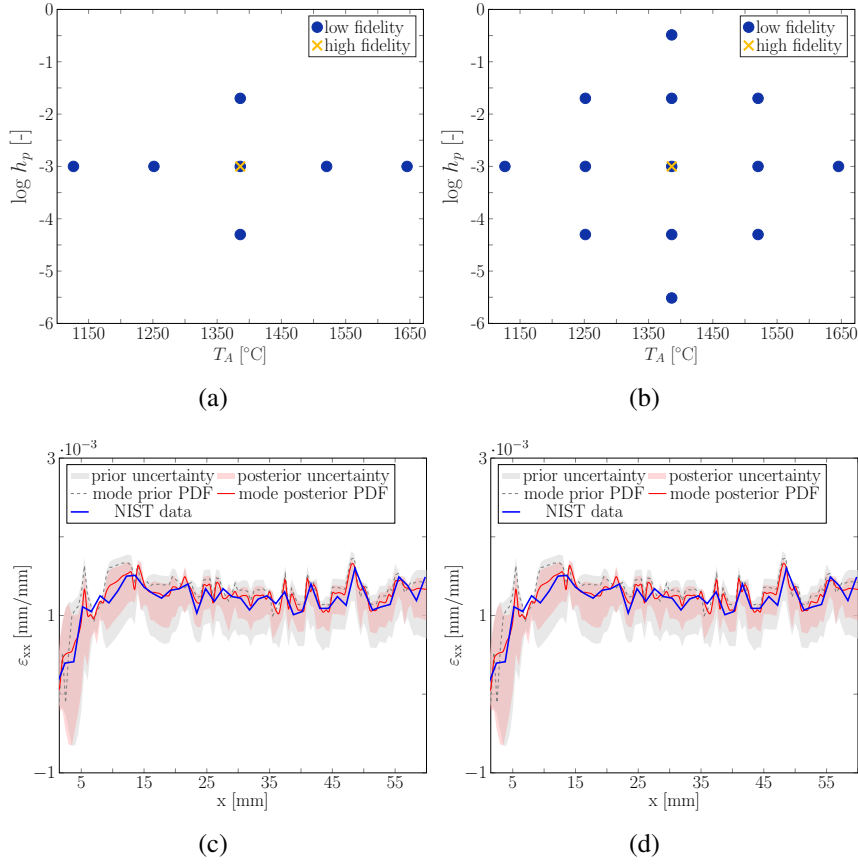


Figure 2.29: Results of the data-informed forward UQ analysis. Top row: collocation points for the MISC surrogate model with 7 (left) and 13 (right) low-fidelity evaluations. Bottom row: corresponding modes and uncertainty bands of the residual strains.

computational advantage of the multi-fidelity approach, which is able to achieve comparable predictive behavior at a substantially reduced computational cost.

2.6 Conclusions

The chapter presented and examined two complementary workflows for the uncertainty quantification in a thermomechanical numerical model of a cantilever beam produced by laser-based powder bed fusion of metals. Both workflows follow the same structured plan, consisting of the identification of the most influential sources of uncertainty, the estimation of the input quantities through comparison with data, and the prediction of residual strains accompanied by an assessment of the related uncertainty. The difference between the two workflows lies in the construction of the surrogate models and in the type of data used for the estimation process.

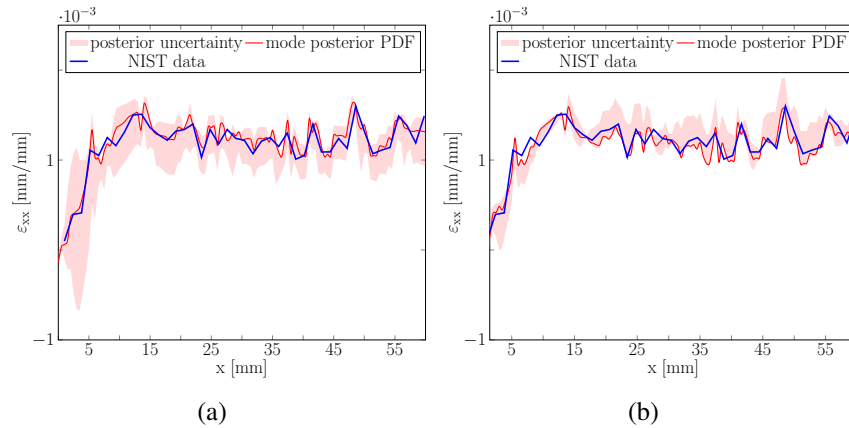


Figure 2.30: Results of the data-informed forward UQ analysis. Comparison in terms of prediction of residual strains of the PBF-LB/M numerical model and associated uncertainty bands using (a) surrogate models based on multi-fidelity PBF-LB/M numerical model evaluations; (b) surrogate models based on high-fidelity PBF-LB/M numerical model evaluations.

The first workflow employed surrogate models constructed exclusively from numerical simulations performed on a single computational mesh. The sensitivity analysis showed that only two of the input quantities significantly influenced the predicted cantilever beam deflection, thus allowing a reduction in the dimension of the analysis. The estimation step was then carried out using synthetic measurements free of measurement noise. This analysis revealed different degrees of identifiability for the two inputs, which motivated the adoption of a mixed probabilistic representation based on a Gaussian distribution for one quantity and a uniform distribution on a reduced interval for the other. Using these data-informed distributions, the prediction of residual strains produced modal profiles in good agreement with the synthetic reference and substantially narrower uncertainty bands in comparison with those obtained from the initial uninformed distributions. The entire workflow, which also included a posteriori validation of the surrogate models, required a limited number of complete numerical simulations, highlighting the efficiency of surrogate modeling based on sparse grids.

The second workflow sought to further reduce computational cost by combining inexpensive numerical simulations performed on a coarse mesh with a small number of simulations on a fine mesh. Surrogate models were constructed through a multi-index approach capable of exploiting the correlation between the two discretizations. In this case, the estimation step was conducted using experimental measurements of the cantilever beam deflection. The most probable values of the uncertain inputs were obtained, together with a Gaussian approximation of their posterior distribution characterized by a covariance matrix with negligible statistical dependence between the two quantities.

It is important to remark that such Gaussian description is obtained via a local Laplace approximation of the posterior distribution around the MAP point and does not preserve the

bounded support of the prior distribution. As a consequence, the approximate posterior may extend beyond the admissible parameter domain, which reflects a limitation of the approximation rather than a violation of Bayes' theorem.

Despite this limitation, the resulting approximation provides a consistent local characterization of the uncertainty and enables efficient uncertainty propagation. The forward analysis shows that the multi-fidelity surrogate models generate residual-strain profiles that closely match the experimental reference and produce uncertainty bands that are, on average, significantly narrower than those obtained from the uninformed prior distributions. A direct comparison with a surrogate model built exclusively from fine-mesh simulations shows comparable predictive behavior, while requiring a substantially higher computational cost.

Overall, the two workflows demonstrate the effectiveness of the adopted methodology. The sensitivity analysis enables a principled reduction of the uncertain-input space; the estimation step provides not only the most probable values of the input quantities but also a probabilistic characterization suitable for propagation; and the prediction of residual strains delivers probabilistic information that accounts for both intrinsic and model-related uncertainties. The single-fidelity workflow shows that surrogate models based on sparse grids make complete uncertainty quantification feasible at moderate cost, whereas the multi-fidelity workflow illustrates how the combination of different numerical discretizations improves the balance between accuracy and computational efficiency.

Several limitations remain. The estimation procedures rely on Gaussian or mixed Gaussian-uniform approximations of the posterior distribution, which are adequate when the likelihood is nearly quadratic but may become inaccurate in the presence of more complex distributions. The modeling of measurement errors is restricted to independent Gaussian noise, and no explicit model-discrepancy correction is included. Furthermore, the analysis is limited to two uncertain inputs. Future developments will include the incorporation of richer experimental datasets, the adoption of derivative-based strategies to refine the description of the uncertainty in the estimated inputs, and the use of more advanced multi-fidelity approaches capable of further reducing the number of fine-mesh simulations required without compromising the accuracy of the uncertainty predictions.

Chapter 3

Uncertainty-Aware Bayesian Inversion and Adaptive Surrogate Modeling for Inverse Design

The present chapter reports two research contributions developed during the doctoral program, both currently under review and not yet published in peer-reviewed venues. The focus is on surrogate-based approaches for inverse problems, combining adaptive surrogate modeling with Bayesian inversion.

Section 3.1 presents a general formulation for inverse problems based on Gaussian-process surrogate models constructed through adaptive sampling strategies, and their integration within a Bayesian inversion framework.

Section 3.2 considers the application of this approach to the inverse design of spinodoid architected materials, where the surrogate model is employed to approximate the structure–property mapping and to support the identification of parameter configurations associated with prescribed effective responses.

The two sections address complementary aspects of surrogate-based Bayesian inversion, encompassing both methodological developments and their application to a representative engineering problem.

3.1 An Uncertainty-Aware Framework for Inverse Problems through Adaptive Surrogate Modeling and Bayesian Inversion

3.1.1 Introduction

Inverse problems [36–38, 127–131] arise in a wide range of scientific and engineering applications and provide a systematic framework for inferring unknown model parameters from observed data. In contrast to forward problems, which aim at predicting quantities of interest from prescribed inputs, inverse problems seek parameter configurations that are consistent with measured or target responses. Applications include structural health monitoring [132], geophysics [133], materials design [134, 135], nanophotonics [126], and fluid-related systems [136, 137], where

direct identification of governing parameters is often impractical or computationally prohibitive.

Inverse problems are typically characterized by ill-posedness, nonlinearity, and non-convexity, which may result in instability, non-uniqueness of solutions, and strong sensitivity to noise and modeling assumptions [36, 37, 129]. Such characteristics motivate formulations that incorporate prior information and explicitly account for uncertainty. Within this context, uncertainty quantification provides a principled framework to assess the impact of measurement noise, limited data, and model inadequacy on parameter inference and predictive reliability [138–140].

A major computational challenge in inverse analysis arises from the repeated evaluation of expensive forward models. High-fidelity numerical simulations provide accurate descriptions of complex physical processes, yet their use within iterative inference procedures may become computationally prohibitive. Surrogate models address this limitation by replacing the original model with computationally efficient approximations [27, 141–143]. Among the available approaches, Gaussian process surrogate models are particularly attractive due to their flexibility in representing nonlinear mappings and their ability to provide a probabilistic characterization of predictive uncertainty [144–147].

The construction of accurate surrogate models under limited evaluation budgets motivates the use of adaptive sampling strategies, whereby new model evaluations are selected based on the current surrogate predictions and associated uncertainty estimates [148–151]. Such approaches enable progressive refinement of the surrogate model over the parameter domain while controlling computational cost [152–155]. In this setting, the primary objective is to obtain a surrogate model that provides a sufficiently accurate approximation of the forward response over the domain of interest.

Bayesian inversion provides a rigorous probabilistic formulation of inverse problems by combining prior information and observational data through Bayes’ theorem [37, 38, 156, 157]. The solution is expressed as a posterior distribution over the unknown parameters, which supports both parameter estimation and uncertainty quantification. In practice, posterior exploration is commonly performed using sampling-based techniques such as Markov chain Monte Carlo (MCMC) [105, 158, 159], which provide asymptotically exact inference at the cost of a large number of model evaluations. To alleviate this computational burden, surrogate-assisted Bayesian inversion approaches have been proposed, in which surrogate models are used to approximate the forward response and accelerate likelihood evaluations [30, 160–164].

Within this context, a practical strategy consists in constructing an accurate surrogate model of the forward problem and subsequently employing it within a Bayesian inversion framework. Such an approach decouples surrogate construction from the inference stage while enabling a substantial reduction in computational cost during posterior evaluation. The accuracy of the resulting parameter inference is therefore inherently linked to both the approximation quality of the surrogate model and the properties of the adopted inference strategy.

The present work adopts this perspective and introduces an uncertainty-aware framework for inverse problems, in which uncertainty is explicitly quantified both at the surrogate modeling level and within the Bayesian inversion formulation. The proposed approach constructs Gaussian process surrogate models through uncertainty-aware adaptive sampling and subsequently employs the resulting surrogate for parameter inference.

A key feature of the framework lies in the explicit treatment of posterior structure in the inference stage, once the surrogate model has been constructed. Parameter estimation is performed through a maximum a posteriori formulation [107, 108], while uncertainty characterization is adapted to the local properties of the posterior distribution. In particular, Laplace approximations [111, 165] are employed in regions where the posterior can be considered locally unimodal, whereas diagnostic indicators based on the negative least-squares functional are used to identify potential multimodality and to localize high-probability regions.

The main contribution of the present work lies in the formulation and systematic assessment of a surrogate-assisted Bayesian framework that combines uncertainty-aware surrogate construction with posterior-based parameter inference. The analysis is conducted on controlled low-dimensional benchmark problems, which enable a detailed investigation of the interplay between surrogate model accuracy, adaptive refinement, and posterior structure.

The numerical study considers one- and two-dimensional analytical benchmark functions, including the Mixed Gaussian–Periodic, Lévy, Griewank, Forrester, and Rosenbrock functions. These benchmarks exhibit nonlinearity, multimodality, and non-convexity, and therefore provide a suitable setting for assessing surrogate model accuracy, parameter estimation behaviour, and uncertainty characterization. The focus on low-dimensional problems allows a controlled and interpretable analysis of the proposed methodology without introducing additional complexities related to high-dimensional settings.

The remainder of the manuscript is organized as follows. Section 3.1.2 reviews the relevant literature. Section 3.1.3 presents the proposed framework. Section 3.1.4 reports the numerical validation. Finally, Section 3.1.5 summarizes the main findings and outlines future developments.

3.1.2 Related Work

3.1.2.1 Surrogate Modeling for Inverse Problems

Surrogate modeling is widely employed to alleviate the computational cost associated with inverse problems, particularly when the underlying forward model is expensive to evaluate [27, 141, 142]. In such settings, surrogate models provide computationally efficient approximations of the forward response and enable the repeated evaluations required for parameter estimation and uncertainty analysis.

A broad range of surrogate modeling techniques has been developed [141, 142], including polynomial-based approximations [166–168], sparse grid methods [76, 77, 80], radial basis function interpolation, and Gaussian process (GP) regression [145, 169, 170]. Among these approaches, GP models are particularly attractive due to their nonparametric formulation and their ability to provide predictive mean and uncertainty estimates in a unified probabilistic setting [144, 146, 147]. Such features make GP surrogates especially suitable for inverse problems, where uncertainty in the forward approximation propagates to parameter inference.

Within Bayesian inversion frameworks, surrogate models have been employed to approximate either the forward model or the likelihood function, thereby reducing the computational burden associated with posterior evaluation [30, 160–164]. These approaches enable the application of

Bayesian methods to computationally demanding problems, while preserving the probabilistic structure of the formulation.

In this context, the construction of surrogate models is typically driven by the objective of achieving sufficiently accurate approximations of the forward response over the parameter domain of interest. The resulting surrogate accuracy plays a central role in determining the reliability of subsequent parameter inference and uncertainty quantification.

3.1.2.2 Adaptive Sampling and Uncertainty-Driven Strategies

Adaptive sampling techniques aim to improve surrogate model accuracy under limited computational budgets by iteratively enriching the training dataset [152, 153, 155]. These approaches exploit information obtained during the surrogate construction process to guide the selection of new evaluation points.

A widely adopted class of methods relies on uncertainty-driven criteria, where sampling decisions are informed by surrogate model uncertainty or error indicators. In the case of Gaussian process regression, the predictive variance provides a natural measure of uncertainty that can be used to identify regions where additional data are expected to improve the surrogate model [147, 152].

Other approaches draw on principles from experimental design and information theory, such as variance reduction and entropy-based criteria, with the objective of improving the overall approximation properties of the surrogate model across the parameter domain. These strategies are generally formulated to enhance global surrogate accuracy while maintaining computational efficiency.

In the context of inverse problems, adaptive sampling has also been investigated as a means to support parameter inference, for instance by focusing evaluations in regions that influence the likelihood function or the posterior distribution [151, 161]. More recent contributions have explored sequential design strategies explicitly tailored to Bayesian inversion [114, 149, 150]. Nevertheless, the development of adaptive sampling strategies that balance global surrogate accuracy with the requirements of parameter inference remains an active area of research.

3.1.2.3 Bayesian Inversion for Parameter Inference

Bayesian inversion provides a probabilistic framework for parameter estimation by combining prior information and observed data through Bayes' theorem [37, 38, 156, 157]. The solution of the inverse problem is expressed as a posterior distribution over the unknown parameters, which characterizes both the most probable parameter values and the associated uncertainty.

The exploration of posterior distributions is commonly performed using sampling-based techniques such as Markov chain Monte Carlo (MCMC), which provide asymptotically exact representations of the posterior distribution [105, 158]. However, these methods may require a large number of model evaluations and can become computationally demanding in complex or high-dimensional settings.

To address this limitation, surrogate-assisted Bayesian inversion has been extensively investigated, where surrogate models are used to approximate the forward model within likelihood evaluations [30, 161, 162]. These approaches significantly reduce computational cost while maintaining a probabilistic formulation, although their accuracy depends on the fidelity of the surrogate model.

Alternative strategies rely on deterministic approximations of the posterior distribution, such as maximum a posteriori (MAP) estimation combined with local approximations based on the Laplace method [111, 112, 165]. These approaches provide computationally efficient estimates of the posterior mode and local uncertainty, while relying on local approximations that may not capture complex posterior structures such as multimodality.

Overall, existing contributions highlight the central role of surrogate modeling and adaptive sampling in enabling efficient Bayesian inversion. Within this context, the present work considers a framework in which surrogate construction and Bayesian inversion are combined in a sequential manner, with the objective of achieving accurate and computationally efficient parameter inference while explicitly accounting for uncertainty.

3.1.3 Proposed Uncertainty-Aware Framework

The present section introduces the proposed uncertainty-aware framework for parameter inference, which combines adaptive surrogate modeling and Bayesian inversion within a unified computational setting. The objective is to identify input parameters that are consistent with a given observed quantity of interest, while providing a probabilistic characterization of the associated epistemic uncertainty.

To alleviate the computational burden associated with repeated evaluations of high-fidelity models, the framework relies on a surrogate approximation of the forward mapping. The surrogate model is constructed through an adaptive sampling strategy driven by predictive uncertainty and is subsequently employed within a Bayesian inversion procedure to approximate the forward model in the likelihood evaluation. This sequential strategy, in which surrogate construction precedes inference, is consistent with widely adopted surrogate-assisted Bayesian approaches [30, 161, 162] and enables a substantial reduction in computational cost while preserving a probabilistic formulation.

The overall workflow consists of three stages — initialization, adaptive surrogate modeling, and Bayesian inversion — as illustrated in Figure 3.1. Starting from an initial dataset of input–output pairs, a GP surrogate model is constructed and progressively refined by selecting new evaluation points in regions characterized by high predictive uncertainty, as quantified by the predictive variance of the surrogate model.

Once a prescribed level of accuracy is achieved, the surrogate model is employed in the inversion stage, where the discrepancy between the observed quantity of interest and the surrogate prediction is quantified through a least-squares functional. Under Gaussian noise assumptions, this functional induces the likelihood, which in turn defines the posterior distribution used for parameter estimation and uncertainty quantification [37].

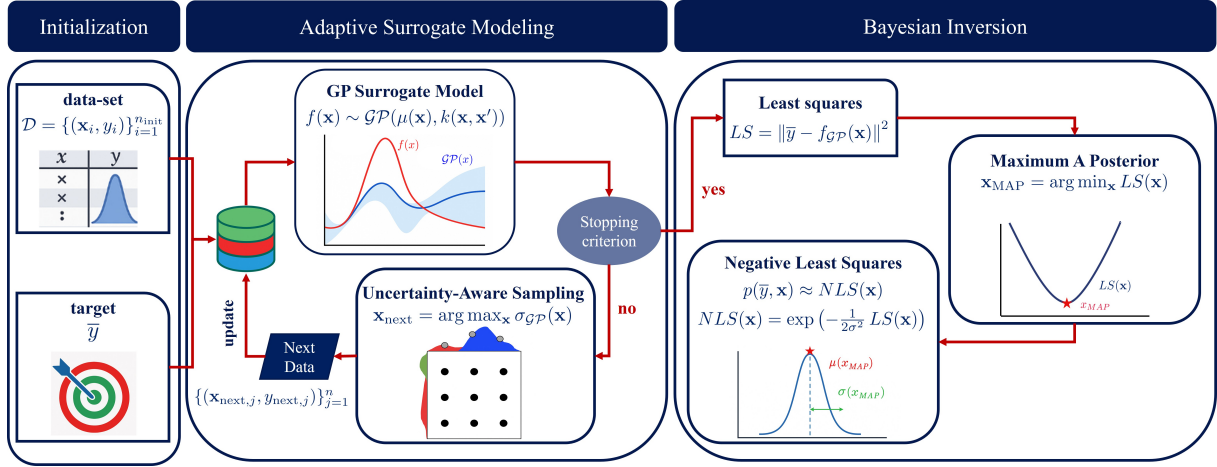


Figure 3.1: Schematic representation of the proposed uncertainty-aware framework for inverse problems. The workflow consists of three stages: initialization, adaptive surrogate modeling, and Bayesian inversion. An initial dataset is used to construct a Gaussian process (GP) surrogate model, which is iteratively refined through uncertainty-aware sampling based on its predictive variance. Once a prescribed accuracy is achieved, the surrogate model is employed within a Bayesian inversion procedure, where the least-squares functional induces the likelihood and enables the estimation of the posterior distribution.

3.1.3.1 Problem Formulation

Let $\mathbf{x} \in \mathbb{R}^d$ denote the vector of input parameters and let $y = f(\mathbf{x}) \in \mathbb{R}$ denote the corresponding quantity of interest obtained from a high-fidelity model $f : \mathbb{R}^d \rightarrow \mathbb{R}$. Given an observed quantity of interest $\bar{y} \in \mathbb{R}$, the inverse problem consists in identifying the parameter values \mathbf{x} that are consistent with the observation, accounting for uncertainty in both data and model representation. In general, this problem is ill-posed in the sense of Hadamard, as existence, uniqueness, or stability of the solution may not be guaranteed [36, 129].

3.1.3.2 Adaptive Surrogate Modeling

The first stage of the proposed framework consists in constructing a surrogate model that approximates the high-fidelity mapping from input parameters to quantities of interest. The process starts from an initial training dataset:

$$D = \{(\mathbf{x}_i, y_i)\}_{i=1}^{n_{\text{init}}}, \quad y_i = f(\mathbf{x}_i), \quad (3.1)$$

where $\mathbf{x}_i \in \mathbb{R}^d$ are input parameters and $y_i \in \mathbb{R}$ are the corresponding model evaluations.

A Gaussian process (GP) surrogate model is employed to approximate the forward mapping. GP regression provides a nonparametric Bayesian approach to function approximation, in which the unknown function is modeled as a random field characterized by a mean function and a covariance kernel [145]. Conditioned on the dataset D , the surrogate defines a predictive

distribution at any input location \mathbf{x} , characterized by a mean function $\mu_{\mathcal{GP}}(\mathbf{x})$ and a variance $\sigma_{\mathcal{GP}}^2(\mathbf{x})$, such that

$$f(\mathbf{x}) \approx f_{\mathcal{GP}}(\mathbf{x}), \quad f_{\mathcal{GP}}(\mathbf{x}) \sim \mathcal{N}(\mu_{\mathcal{GP}}(\mathbf{x}), \sigma_{\mathcal{GP}}^2(\mathbf{x})), \quad (3.2)$$

where the predictive moments are induced by the GP prior defined through a mean function $\mu(\mathbf{x})$ and a covariance kernel $k(\mathbf{x}, \mathbf{x}')$.

The surrogate model is iteratively refined through an adaptive sampling strategy driven by predictive uncertainty. At each iteration, new sampling points are selected according to

$$\mathbf{x}_{\text{next}} = \arg \max_{\mathbf{x} \in \mathcal{X}} \sigma_{\mathcal{GP}}(\mathbf{x}), \quad (3.3)$$

which promotes exploration of regions characterized by high predictive uncertainty. This choice is consistent with classical uncertainty-based sampling criteria in GP modeling, where the predictive variance is used as an indicator of model uncertainty [142, 152].

The dataset is then enriched as

$$D \leftarrow D \cup \{(\mathbf{x}_{\text{next},j}, y_{\text{next},j})\}_{j=1}^n, \quad y_{\text{next},j} = f(\mathbf{x}_{\text{next},j}), \quad (3.4)$$

and the surrogate model is retrained on the updated dataset.

The adaptive procedure is repeated until a prescribed convergence criterion is satisfied. In the present study, convergence is assessed through the mean squared error (MSE) evaluated on an independent validation set of samples n_{val} :

$$\text{MSE} = \frac{1}{n_{\text{val}}} \sum_{i=1}^{n_{\text{val}}} (y_i - \mu_{\mathcal{GP}}(\mathbf{x}_i))^2, \quad (3.5)$$

where $y_i = f(\mathbf{x}_i)$ are high-fidelity evaluations.

The use of the MSE as a convergence criterion is justified in controlled benchmark settings, where access to the high-fidelity model enables a direct and quantitative assessment of surrogate accuracy. In practical applications, where such evaluations are expensive, alternative stopping criteria may be employed, including the stabilization of acquisition functions, reduction of predictive variance, or marginal improvement of surrogate predictions [171]. A systematic comparison of these strategies is beyond the scope of the present work.

3.1.3.3 Bayesian Inversion Framework

Once the surrogate model achieves the desired level of accuracy, parameter inference is performed through Bayesian inversion. Given an observed quantity of interest $\bar{y} \in \mathbb{R}$, the objective is to infer the parameter vector $\mathbf{x} \in \mathbb{R}^d$ while quantifying the associated uncertainty.

Bayes' theorem defines the posterior distribution as

$$p(\mathbf{x} | \bar{y}) \propto p(\bar{y} | \mathbf{x}) p(\mathbf{x}), \quad (3.6)$$

where $p(\bar{y} | \mathbf{x})$ is the likelihood function and $p(\mathbf{x})$ is the prior distribution [37].

The likelihood is evaluated using the surrogate model $f_{\mathcal{GP}}(\mathbf{x})$, which replaces the high-fidelity model. Assuming additive Gaussian noise, the observation model is given by

$$\bar{y} = f_{\mathcal{GP}}(\mathbf{x}) + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 I). \quad (3.7)$$

Under this assumption, the likelihood takes the form

$$p(\bar{y} | \mathbf{x}) \propto \exp\left(-\frac{1}{2\sigma^2} \|\bar{y} - f_{\mathcal{GP}}(\mathbf{x})\|^2\right), \quad (3.8)$$

which corresponds to a Gaussian likelihood induced by a least-squares misfit functional.

A commonly adopted and analytically convenient choice for the prior distribution is a multivariate Gaussian [107]:

$$p(\mathbf{x}) \propto \exp\left(-\frac{1}{2}(\mathbf{x} - \mu_{\mathbf{x}})^\top \Gamma^{-1}(\mathbf{x} - \mu_{\mathbf{x}})\right), \quad (3.9)$$

which encodes prior information on the parameter values.

The MAP estimate is obtained as

$$\mathbf{x}_{\text{MAP}} = \arg \min_{\mathbf{x}} \left[\frac{1}{2\sigma^2} \|\bar{y} - f_{\mathcal{GP}}(\mathbf{x})\|^2 + \frac{1}{2}(\mathbf{x} - \mu_{\mathbf{x}})^\top \Gamma^{-1}(\mathbf{x} - \mu_{\mathbf{x}}) \right], \quad (3.10)$$

which corresponds to the minimization of the negative log-posterior.

In the absence of informative prior knowledge, a uniform prior leads to the least-squares formulation

$$\mathbf{x}_{\text{MAP}} = \arg \min_{\mathbf{x}} \|\bar{y} - f_{\mathcal{GP}}(\mathbf{x})\|^2, \quad (3.11)$$

which highlights the connection between Bayesian inversion and deterministic inverse formulations.

When the posterior distribution is locally unimodal and sufficiently regular, a Laplace approximation centered at the MAP estimate can be employed [111, 165]:

$$p(\mathbf{x} | \bar{y}) \approx \mathcal{N}(\mathbf{x}_{\text{MAP}}, \Sigma_{\text{MAP}}), \quad (3.12)$$

where the covariance Σ_{MAP} is given by the inverse Hessian of the negative log-posterior evaluated at \mathbf{x}_{MAP} . This approximation provides a computationally efficient characterization of local posterior uncertainty, while its validity depends on the regularity and local Gaussianity of the posterior distribution.

3.1.4 Numerical Validation on Analytical Benchmarks

The present section validates the proposed uncertainty-aware framework for inverse problems, introduced in Section 3.1.3, through a series of one- and two-dimensional analytical benchmark

functions. These analytical benchmarks provide a controlled setting for quantitatively evaluating the performance of the framework in terms of surrogate modeling accuracy, robustness of parameter inference, computational efficiency, and uncertainty quantification.

The section is organized as follows. Section 3.1.4.1 presents the validation results for one-dimensional benchmarks, analysing separately the construction of the surrogate model (Section 3.1.4.1.1), parameter inference (Section 3.1.4.1.2), and a concluding discussion of the results (Section 3.1.4.1.3). Section 3.1.4.2 extends the validation to two-dimensional benchmarks, following the same structure to ensure consistency between surrogate modeling (Section 3.1.4.2.1), parameter inference (Section 3.1.4.1.2), and a concluding discussion of the results (Section 3.1.4.2.3). Each case includes an assessment of uncertainty quantification and computational cost, which are critical aspects in practical inverse problem scenarios.

All numerical experiments are conducted using the open-source Python library UQforPy [172], which provides a modular and extensible platform for uncertainty quantification and inverse design. Additional implementation details are provided in C.

3.1.4.1 One-Dimensional Benchmarks

The present section validates the proposed framework (Section 3.1.3) using the one-dimensional analytical benchmark functions illustrated in Figure 3.2. The selected benchmarks capture diverse levels of functional complexity, including strong nonlinearity, multimodality, and high-frequency oscillations, and are widely adopted to assess the performance of surrogate modeling strategies and parameter inference in inverse problems [37, 157, 162, 173].

The benchmark suite includes four representative functions:

- *Mixed Gaussian-Periodic Function:*

$$y = \exp\left(-\frac{(x-2)^2}{2}\right) + 0.9 \exp\left(-\frac{(x+5)^2}{20}\right) - 0.1 \cos(2x) \quad \text{with } x \in [-10, 10] \quad (3.13)$$

combining localized Gaussian peaks with periodic oscillations, resulting in both global smoothness and local high-frequency variations (Figure 3.2(a)).

- *Lévy Function:*

$$y = (\sin(\pi w))^2 + (w-1)^2 [1 + 10(\sin(\pi w + 1))^2] + (w-1)^4 (\sin(2\pi w))^2 \quad (3.14)$$

where $w = 1 + \frac{x-1}{4}$ and $x \in [-6, 6]$. The Lévy function presents a rugged, multimodal landscape with many local minima (Figure 3.2(b)), posing challenges to both global exploration and surrogate model fidelity.

- *Griewank Function:*

$$y = \frac{x^2}{4000} - \cos(x) + 1 \quad \text{with } x \in [-15, 15] \quad (3.15)$$

combining a slow quadratic trend with fine-scale periodic modulation (Figure 3.2(c)).

- *Forrester Function*:

$$y = (6x - 2)^2 \sin(12x - 4) \quad \text{with } x \in [0, 1] \quad (3.16)$$

characterized by smooth global behaviour with sharply localized nonlinearities (Figure 3.2(d)).




Figure 3.2: One-dimensional analytical benchmark functions used to validate the proposed uncertainty-aware framework for inverse problems. The suite includes representative challenges such as smooth nonlinearity, multimodal structure, and fine-scale variability, allowing systematic assessment of surrogate modeling and parameter inference capabilities. (a) Mixed Gaussian-Periodic. (b) Lévy. (c) Griewank. (d) Forrester.

3.1.4.1.1 Adaptive Surrogate Model Construction

The present section assesses the surrogate modeling capabilities of the proposed (Section 3.1.3) on the one-dimensional analytical benchmark functions introduced in Figure 3.2, with particular focus on the Mixed Gaussian-Periodic case (Figure 3.2(a)). Since the surrogate modeling procedure remains identical across all benchmark functions (Figure 3.2(b)–(d)), the analysis is primarily conducted on the Mixed Gaussian-Periodic function, while additional results for the remaining benchmarks are reported in D. Final surrogate models for all benchmark functions are summarized at the end of the present section.

Surrogate models are constructed using GP surrogate models combined with an adaptive uncertainty-aware sampling strategy. At each iteration, new sampling points are selected by maximizing the predictive standard deviation of the surrogate model, consistently with the adaptive procedure introduced in Section 3.1.3.2. This strategy promotes exploration of regions characterized by high epistemic uncertainty, leading to a progressive improvement of the surrogate model over the entire input domain.

The performance of the surrogate models is evaluated in terms of both predictive accuracy and sampling efficiency. The proposed GP-based surrogate modeling approach is compared against classical deterministic alternatives, including Lagrangian Polynomial (LP) interpolation, Legendre Expansion (LE), and Cubic Spline (CS) interpolation [166–168, 174, 175]. Two covariance kernels are considered for the GP surrogate models: the Radial Basis Function

(RBF) kernel [176] with length scale $\ell = 1.0$, and the Matérn kernel [170] with smoothness parameter $\nu = 5/2$ and identical length scale. All surrogate models are trained on datasets generated from the high-fidelity model, and their predictive accuracy is quantified using the MSE (Equation (3.5)) computed over an independent validation set of $n_{\text{val}} = 1000$ samples, ensuring statistical robustness of the reported metrics.

Figure 3.3 illustrates the comparative performance of the surrogate models. As shown in Figure 3.3(a), GP surrogate models accurately capture the multimodal and fine-scale structure of the Mixed Gaussian–Periodic benchmark function, with the Matérn kernel exhibiting enhanced robustness in highly oscillatory regions due to its increased flexibility in modeling non-smooth behaviour. Among deterministic alternatives, CS interpolation performs reasonably well in smooth regions but lacks the ability to resolve localized nonlinearities. LP interpolation exhibits Runge-type instabilities at the domain boundaries [177], while LE maintains numerical stability but suffers from limited representational capacity. Quantitative comparison via MSE (Figure 3.3(b)) confirms the superior accuracy of GP surrogate models, with the Matérn kernel consistently achieving the lowest prediction error. These observations are further supported by the convergence analysis reported in Figure 3.3(c), which highlights the improved sample efficiency of GP-based surrogate modeling.

A key feature of GP surrogate models is the availability of predictive uncertainty estimates alongside mean predictions. This enables the construction of adaptive sampling strategies driven by epistemic uncertainty, as formalized in Equations (3.3) and (3.4), leading to a targeted enrichment of the training dataset.

Figure 3.4 illustrates the progressive refinement of the GP surrogate model for the Mixed Gaussian–Periodic benchmark function. Starting from a sparse initial dataset of 5 samples, the surrogate model is iteratively improved by acquiring new samples in regions characterized by high predictive uncertainty. This results in a rapid reduction of uncertainty and a significant improvement in predictive accuracy, particularly in regions exhibiting localized variations. The observed behaviour highlights the effectiveness of uncertainty-aware sampling in constructing accurate surrogate models under limited data regimes.

Final surrogate models obtained for all benchmark functions are shown in Figure 3.5. The configuration settings used across all benchmark functions are reported in Table 3.1. Despite the diversity in functional complexity — ranging from multimodal landscapes (Mixed Gaussian–Periodic, Lévy), to oscillatory behaviour (Griewank), and localized nonlinearities (Forrester) — the proposed framework consistently delivers accurate surrogate models using a limited number of samples within the considered one-dimensional setting.

The results confirm the accuracy, robustness, and sample efficiency of the proposed surrogate modeling strategy, establishing a reliable foundation for the subsequent parameter inference tasks in low-dimensional benchmark problems. GP surrogate models consistently outperform classical deterministic alternatives, particularly in low-data regimes and in the presence of sharp nonlinearities or oscillatory structures.

It is worth emphasizing that the adopted adaptive sampling strategy is designed to promote a balanced exploration of the input space by targeting regions of high epistemic uncertainty. This choice reflects the objective of constructing a globally accurate surrogate model that faithfully

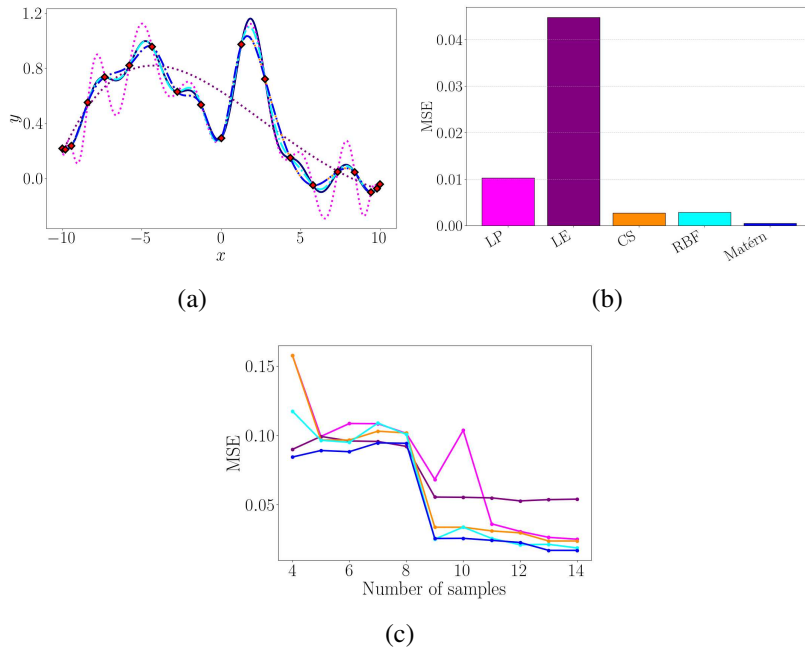


Figure 3.3: Performance comparison of surrogate models for the Mixed Gaussian–Periodic benchmark function. (a) Predictive responses obtained with different surrogate models. The solid black line denotes the high-fidelity model; dotted lines correspond to LP (fuchsia), LE (purple), and CS (orange) deterministic alternatives; cyan and blue denote GP-based surrogate models with Matérn and RBF kernels, respectively. Red diamonds indicate number of samples. (b) MSE values for 14 number of samples. (c) MSE evolution with increasing number of samples.

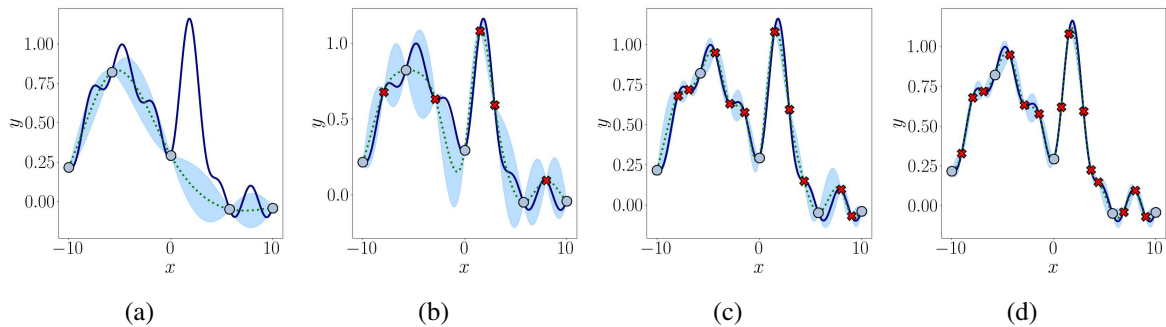


Figure 3.4: Progressive refinement of the GP surrogate model for the Mixed Gaussian–Periodic benchmark function using the uncertainty-aware sampling. Blue solid line: high-fidelity model; green dotted line: GP mean (Matérn); light blue band: 95% confidence interval. Gray circles: initial samples; red crosses: acquired new samples. (a) 5 number of samples. (b) 10 number of samples. (c) 15 number of samples. (d) 19 number of samples.

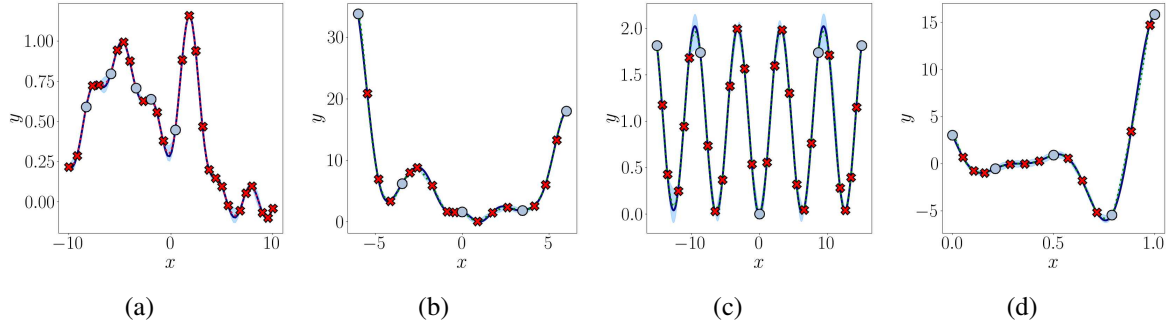


Figure 3.5: Final surrogate models for all benchmark functions. Blue solid line: high-fidelity; red dotted line: GP (Matérn); light blue band: predictive uncertainty. Gray circles: initial samples; red crosses: acquired new samples. (a) Mixed Gaussian–Periodic. (b) Lévy. (c) Griewank. (d) Forrester.

Table 3.1: Configuration settings for surrogate modeling across benchmark functions.

Parameter	Value
Initial samples (n_{init})	5
Acquired new samples (n_{acq})	15–20
Surrogate model	Gaussian Process (GP)
Covariance kernel	Matérn $\nu = 5/2$
Mean function	Zero mean prior
Noise model	Gaussian noise ($\sigma^2 = 10^{-6}$)
Hyperparameters optimization	Log-marginal likelihood maximization
Sampling strategy	Predictive variance maximization
Stopping criterion	Validation MSE $< 10^{-3}$
Validation set (n_{val})	1000 samples

reproduces the high-fidelity response across the parameter domain. While alternative strategies may focus on localized refinement, the present approach prioritizes global model fidelity, which is essential for robust and reliable uncertainty-aware inverse analysis.

3.1.4.1.2 Bayesian Inversion Results

The present section evaluates the capability of the proposed uncertainty-aware framework to perform parameter inference and quantify epistemic uncertainty through Bayesian inversion (Section 3.1.3.3) on the one-dimensional analytical benchmark functions illustrated in Figure 3.2.

The analysis focuses on the Mixed Gaussian–Periodic (Figure 3.2(a)) and Forrester (Figure 3.2(d)) benchmark functions, which highlight distinct inverse problem characteristics. Additional results for the remaining benchmarks are reported in E.

Figure 3.6 presents the results for the Mixed Gaussian–Periodic benchmark function. Figure 3.6(a) shows the optimization outcomes obtained from multiple initial guesses (red triangles) distributed across the parameter domain to ensure robust exploration of the parameter landscape. The optimization is performed using the L-BFGS-B algorithm with box constraints, with a maximum of 400 iterations per run. The observed quantity of interest $\bar{y} = 0.63$ is indicated by the black horizontal line, while the green stars denote the corresponding MAP estimates.




Figure 3.6: Parameter inference via Bayesian inversion for the Mixed Gaussian–Periodic benchmark function. (a) MAP estimates obtained from multiple optimization runs. (b) LS profile across the parameter space. (c) NLS profile representing the posterior distribution approximation and highlighting the high-probability region.

The results in Figure 3.6a reveal the presence of multiple distinct MAP estimates, each corresponding to a different parameter configuration that reproduces the observed quantity of interest. This multiplicity reflects the ill-posed nature of the inverse problem and is directly linked to the multimodal structure of the LS functional, as shown in Figure 3.6(b). The oscillatory behaviour of the forward model induces several local minima, each representing a feasible solution.

The structure of the posterior distribution is further illustrated through the NLS profile in Figure 3.6c. The presence of multiple peaks confirms that the posterior distribution is inherently multimodal, with probability mass concentrated around distinct regions of the parameter space. In such cases, local Gaussian approximations centered at a single MAP estimate may be inadequate to describe the full posterior structure, as they neglect the contribution of secondary modes.

Despite this complexity, the NLS profile provides a useful approximation for identifying high-probability regions. By introducing a threshold level $NLS(x) = 0.95$, the parameter domain can be restricted to a subset containing the most plausible solutions. As shown in Figure 3.6c, the admissible region is reduced from the original interval $[-10, 10]$ to approximately $[-7.5, 5]$, demonstrating the capability of Bayesian inversion to significantly constrain the parameter space.

Figure 3.7 presents the results for the Forrester benchmark function. In contrast to the previous case, the LS profile (Figure 3.7(b)) exhibits a single well-defined minimum, leading to a unique MAP estimate, as shown in Figure 3.7(a). The corresponding NLS profile (Figure 3.7(c)) is sharply peaked and approximately Gaussian, indicating a well-posed inverse problem with low epistemic uncertainty.




Figure 3.7: Parameter inference via Bayesian inversion for the Forrester benchmark function. (a) MAP estimate obtained from multiple optimization runs. (b) LS profile across the parameter space. (c) NLS profile representing the posterior distribution approximation and highlighting the high-probability region.

The posterior distribution exhibits strong concentration around the inferred parameter value $x \approx 0.76$, with a narrow spread of approximately 0.02. This behaviour indicates a high level of confidence in the inferred solution and confirms the identifiability of the inverse problem. The corresponding high-probability region is significantly reduced from the initial domain $[0, 1]$ to approximately $[0.735, 0.78]$, demonstrating the effectiveness of Bayesian inversion in localizing the solution.

Overall, the results highlight the ability of the proposed framework to capture both multimodal and unimodal inverse problem structures. In the presence of non-uniqueness, the methodology provides a comprehensive characterization of the posterior distribution, while in well-posed settings it yields precise and sharply localized parameter estimates. These findings confirm the robustness of the surrogate-assisted Bayesian inversion approach for uncertainty-aware parameter inference in low-dimensional benchmark problems.

3.1.4.1.3 Discussion

The numerical validation results presented in Sections 3.1.4.1.1 and 3.1.4.1.2 demonstrate the effectiveness of the proposed uncertainty-aware framework for inverse problems in one-dimensional benchmark settings. The framework combines adaptive surrogate model construction with Bayesian inversion, enabling both accurate approximation of the forward model and reliable parameter inference with quantified epistemic uncertainty. A complementary discussion on alternative inference strategies is provided in F.

The surrogate models constructed through uncertainty-aware adaptive sampling accurately reproduce the key characteristics of all benchmark functions, including multimodality (Mixed Gaussian–Periodic, Lévy), high-frequency oscillations (Griewank), and localized nonlinearities (Forrester), while requiring a limited number of high-fidelity model evaluations. For instance, in the Mixed Gaussian–Periodic benchmark function, the Gaussian process surrogate model with Matérn kernel captures the complex response of the high-fidelity model with fewer than 20 samples (Figure 3.5(a)). Similarly, for the Forrester function, high predictive accuracy is

achieved with approximately 15 samples (Figure 3.5(d)). These results highlight the sample efficiency of the adopted adaptive sampling strategy, which systematically targets regions of high epistemic uncertainty.

Once constructed, the surrogate model enables an efficient application of Bayesian inversion. The least-squares functional used to identify the MAP estimate is evaluated on the surrogate model, thereby avoiding repeated evaluations of the high-fidelity model. This surrogate-assisted formulation significantly reduces the computational cost of the inverse problem, while preserving the relevant structure of the forward response.

The results also illustrate the capability of the framework to handle both ill-posed and well-posed inverse problems. In the Mixed Gaussian–Periodic benchmark (Figure 3.6), the least-squares functional exhibits multiple local minima, leading to several admissible MAP estimates. This behaviour reflects the intrinsic non-uniqueness of the inverse problem and is consistent with the multimodal structure of the posterior distribution. Although a single optimal parameter cannot be uniquely identified, the analysis of the NLS profile allows the identification of high-probability regions, effectively restricting the admissible parameter space to a subset of plausible solutions. This provides a meaningful characterization of epistemic uncertainty, which cannot be captured by point estimates alone.

In contrast, the Forrester benchmark function represents a well-posed inverse problem with a unique and identifiable solution. The least-squares functional exhibits a single dominant minimum (Figure 3.7(b)), and the corresponding posterior distribution is sharply concentrated around the MAP estimate (Figure 3.7(c)). This behaviour indicates low epistemic uncertainty and confirms the reliability of the inferred parameter. The resulting reduction of the admissible parameter space further demonstrates the effectiveness of Bayesian inversion in extracting informative solutions from observed data.

From a computational perspective, the separation between surrogate model construction and Bayesian inversion yields a substantial efficiency gain. Direct application of Bayesian inversion on the high-fidelity model would require a large number of model evaluations for each optimization run, especially in the presence of multimodality. In contrast, the use of a surrogate model reduces the cost of each evaluation to negligible levels, enabling extensive exploration of the parameter space and efficient uncertainty quantification.

Overall, the results confirm that the proposed framework provides a robust and computationally efficient approach for surrogate-assisted Bayesian inversion. The combination of uncertainty-aware surrogate modeling and probabilistic parameter inference enables accurate reconstruction of the forward model and reliable characterization of epistemic uncertainty, even in challenging inverse problem scenarios.

3.1.4.2 Two-Dimensional Benchmarks

The proposed uncertainty-aware framework for inverse problems, introduced in Section 3.1.3, is further validated through two-dimensional analytical benchmark functions, shown in Figure 3.8. These benchmarks are employed to investigate the framework’s behaviour in the presence of coupled parameters, increased problem complexity with respect to the one-dimensional setting,

and non-trivial functional landscapes, while enabling a controlled and interpretable assessment of surrogate modeling and parameter inference under uncertainty.

The selected benchmark functions are:

- *Mixed Gaussian–Periodic function*, defined over the domain $x \in [-1, 2]$, $y \in [0, 3]$, as

$$f(x, y) = \exp\left(-\frac{(x-2)^2 + (y-2)^2}{2}\right) + 0.2 \cos(3x) \sin(3y) + 0.1 \sin(5x + 5y) \quad (3.17)$$

combining a smooth Gaussian peak with periodic perturbations. This configuration introduces both multimodal behavior and localized nonlinearities, posing challenges to surrogate model construction and parameter inference accuracy. The high-fidelity model is illustrated in Figure 3.8(a) and Figure 3.8(b).

- *Rosenbrock function*, defined over the domain $x \in [-2, 2]$, $y \in [-1, 2]$, as

$$f(x, y) = (1 - x)^2 + 100(y - x^2)^2 \quad (3.18)$$

characterized by a narrow, curved valley and steep gradients. This benchmark is representative of ill-conditioned inverse problems with strong anisotropies and nonlinearity, challenging both the surrogate modeling phase and the parameter inference analysis. The associated high-fidelity response is shown in Figure 3.8(c) and Figure 3.8(d).

3.1.4.2.1 Adaptive Surrogate Model Construction

The present section reports the surrogate modeling results obtained for the two-dimensional benchmark functions introduced in Figure 3.8. The two-dimensional setting is considered to analyse the effect of coupled parameters and increased functional complexity with respect to the one-dimensional case, rather than to assess scalability to higher-dimensional parameter spaces.

Consistently with the methodology described in Section 3.1.3.2 and applied in the one-dimensional setting (Section 3.1.4.1.1), surrogate models are constructed using Gaussian process (GP) surrogate models with a Matérn 5/2 covariance kernel. The observation noise is assumed Gaussian with variance $\sigma^2 = 10^{-6}$. The training dataset is iteratively enriched through an uncertainty-aware adaptive sampling strategy, where new samples are selected by maximizing the predictive standard deviation of the surrogate model, as defined in Equation (3.4). This approach promotes exploration of regions characterized by high epistemic uncertainty, enabling an efficient reconstruction of the response surface.

The analysis is first conducted on the Mixed Gaussian–Periodic benchmark function (Figure 3.8a–b), followed by the Rosenbrock benchmark function (Figure 3.8c–d).

Figure 3.9 reports the surrogate modeling results for the Mixed Gaussian–Periodic benchmark function. The GP surrogate model shown in Figure 3.9a accurately reconstructs the response surface of the high-fidelity model (Figure 3.8(a) and Figure 3.8(b)), capturing both the smooth global Gaussian peak and the superimposed periodic features. The corresponding contour plot in Figure 3.9b illustrates the spatial distribution of the training samples (black markers). Starting from an initial design of 5 samples, additional samples are progressively selected in regions




Figure 3.8: Two-dimensional analytical benchmark functions used to validate the proposed uncertainty-aware framework for inverse problems. (a)–(b): Mixed Gaussian-Periodic function — high-fidelity model surface and contour plots. (c)–(d): Rosenbrock function — high-fidelity model surface and contour plots.

of high predictive uncertainty. The adaptive sampling strategy concentrates samples in areas characterized by strong nonlinearities, particularly around the central peak and regions with oscillatory behaviour, enabling accurate resolution of both global and local features.

The convergence behaviour is quantified in Figure 3.9c, where the MSE (Equation (3.5)) is evaluated on an independent validation set of $n_{\text{val}} = 1000$ samples. A rapid decrease in the error is observed, with convergence effectively achieved after approximately 21 high-fidelity evaluations. This result highlights the sample efficiency of the proposed adaptive surrogate modeling strategy in the presence of coupled parameters and multimodal response surfaces.

Figure 3.10 presents the surrogate modeling results for the Rosenbrock benchmark function. The GP surrogate model in Figure 3.10a successfully reproduces the characteristic narrow and curved valley of the high-fidelity model (Figure 3.8(c) and Figure 3.8(d)), as well as the steep

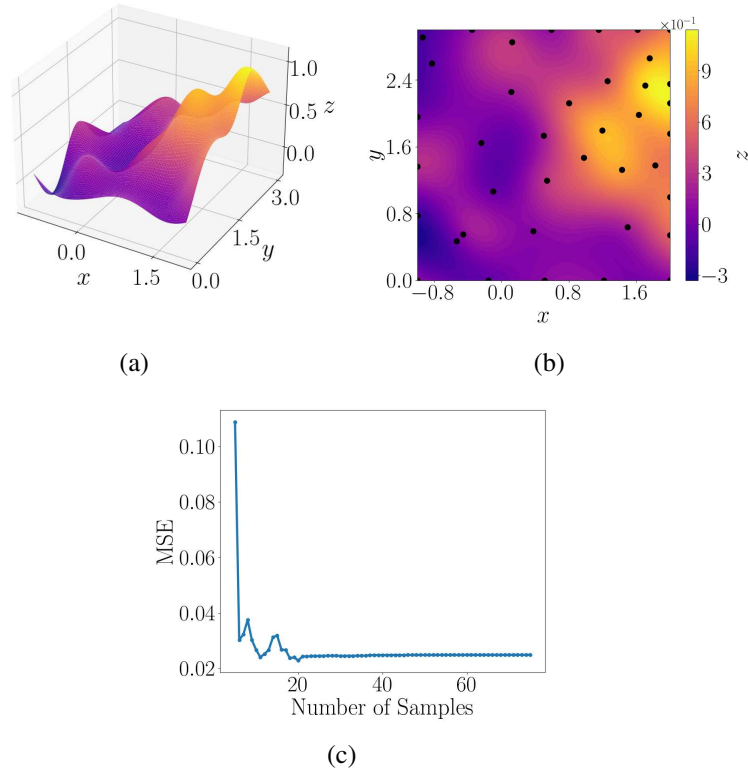


Figure 3.9: Surrogate modeling results for the Mixed Gaussian–Periodic benchmark function. (a) Final surrogate model surface after 21 high-fidelity evaluations. (b) Contour plot with sample locations. (c) Evolution of the mean squared error (MSE) as a function of the number of samples.

gradients in the surrounding regions. The distribution of samples shown in Figure 3.10b reflects the adaptive nature of the sampling strategy: starting from 5 initial samples, additional evaluations are concentrated along the valley, where predictive uncertainty is highest. This targeted sampling enables accurate reconstruction of the strongly anisotropic response surface.

The MSE evolution reported in Figure 3.10c shows a rapid decrease followed by early convergence of the surrogate model accuracy. The stopping criterion is satisfied after approximately 9 high-fidelity evaluations, indicating that the proposed approach efficiently captures the essential features of the Rosenbrock function despite its strong anisotropy and curvature.

3.1.4.2.2 Bayesian Inversion Results

The present section reports the results of parameter inference for the two-dimensional analytical benchmark functions illustrated in Figure 3.8, using the Bayesian inversion framework described in Section 3.1.3.3. The analysis focuses on the two-dimensional setting in order to investigate the behaviour of the inference procedure in the presence of coupled parameters, non-convex response landscapes, and anisotropic structures, rather than to assess scalability to higher-dimensional

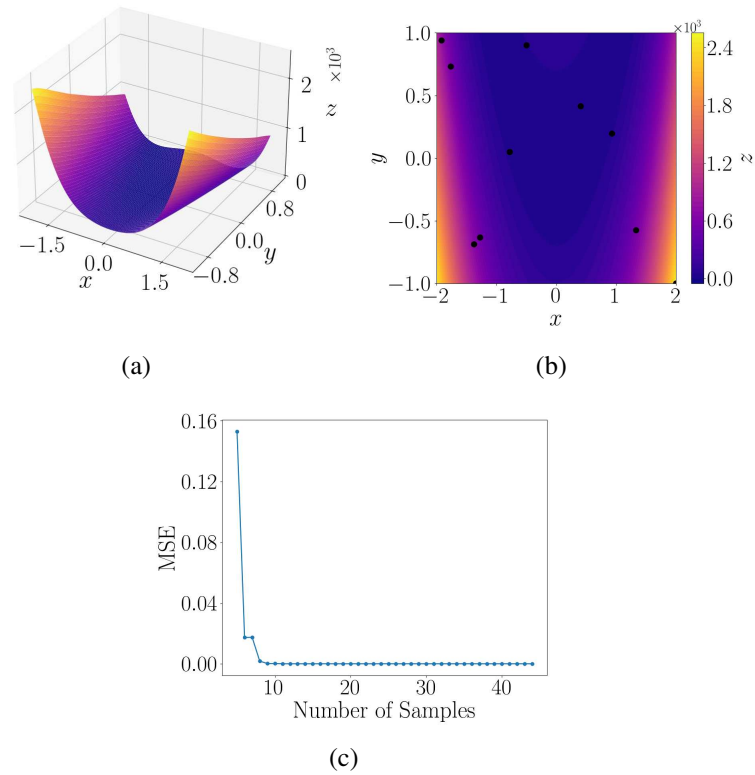


Figure 3.10: Surrogate modeling results for the Rosenbrock benchmark function. (a) Final surrogate model surface after 9 high-fidelity evaluations. (b) Contour plot with sample locations. (c) Evolution of the mean squared error (MSE) as a function of the number of samples.

parameter spaces.

The inverse problems are solved by leveraging the GP surrogate models constructed in Section 3.1.4.2.1 (see Figures 3.9 and 3.10). For each benchmark function, the analysis includes the computation of the MAP estimate, the evaluation of the LS functional, and the corresponding NLS profile, which serves as an approximation of the unnormalized posterior distribution. This enables a systematic characterization of the admissible parameter space and of the associated epistemic uncertainty.

The results are discussed separately for the Mixed Gaussian–Periodic function (Equation (3.17)) and the Rosenbrock function (Equation (3.18)).

Figure 3.11 presents the results for the Mixed Gaussian–Periodic benchmark function. The LS surface shown in Figure 3.11(a) exhibits a non-convex landscape with a dominant basin of attraction, while the corresponding contour plot in Figure 3.11(b) highlights a well-defined low-LS region surrounded by nearly concentric level sets. To ensure adequate exploration of the parameter space, the LS minimization is initialized from multiple independent starting points (red

triangles). Despite the presence of mild non-convexities, the optimization trajectories consistently converge toward a compact region near the prescribed observation $\bar{y} = 0.63$ (black square), yielding multiple closely clustered MAP estimates (green stars) within the same basin.

The inferred solutions concentrate around the parameter vector $\mathbf{x}_{\text{MAP}} = [1.248, 1.812]$, with a residual value $\text{LS} \approx 0.03$, indicating excellent agreement between the surrogate model prediction and the target observation. The corresponding NLS surface and contour plots (Figure 3.11(c), Figure 3.11(d)) reveal a sharply peaked posterior distribution centered at the MAP estimate. The maximum value of the NLS is approximately 0.90, with a rapid decay away from the peak. Although the LS landscape is non-convex, the resulting posterior approximation remains effectively unimodal and strongly localized, indicating low epistemic uncertainty and strong local identifiability.

Parameter-wise uncertainty is quantified through marginal credible intervals derived from a local Gaussian approximation of the posterior distribution around the MAP estimate. The covariance matrix Σ_{MAP} is obtained as the inverse of the Hessian of the negative log-posterior evaluated at the MAP point, as defined in Section 3.1.3.3. The resulting axis-aligned confidence region provides interpretable uncertainty bounds for each parameter. The MAP estimate and associated credible intervals are summarized in Table 3.2.

Table 3.2: MAP estimate and marginal credible intervals for the Mixed Gaussian–Periodic benchmark function. The target value is $\bar{y} = 0.63$, the LS residual at the MAP point is approximately 0.03, and the posterior peak reaches $\text{NLS}_{\text{max}} \approx 0.90$.

Parameter	MAP estimate	Credible interval
x	1.248	[1.158, 1.638]
y	1.812	[1.323, 2.221]

Figure 3.12 presents the results for the Rosenbrock benchmark function. The LS surface (Figure 3.12(a)) captures the characteristic narrow curved valley of the Rosenbrock landscape, while the contour plot (Figure 3.12(b)) confirms the presence of a single dominant minimum. Independent optimization runs initialized from multiple starting points (red triangles) consistently converge to this minimum, resulting in a unique MAP estimate (green star) located near $\mathbf{x}_{\text{MAP}} = [-1.5, -0.6]$, in close proximity to the prescribed observation (black square).

The corresponding NLS surface (Figure 3.12(c), Figure 3.12(d)) is sharply unimodal, with the posterior distribution strongly concentrated around the MAP estimate. No secondary modes are observed, indicating strong local identifiability and low epistemic uncertainty. The geometry of the posterior reflects the anisotropic structure of the underlying forward model, with the uncertainty primarily aligned along the valley direction.

The surrogate models obtained for the two benchmark functions provide sufficiently accurate approximations of the forward response to support Bayesian inversion. This aspect is essential within the proposed framework, since the quality of parameter inference depends directly on the quality of the surrogate approximation in the regions that govern the least-squares functional and, consequently, the posterior distribution. The rapid decrease of the validation error observed in

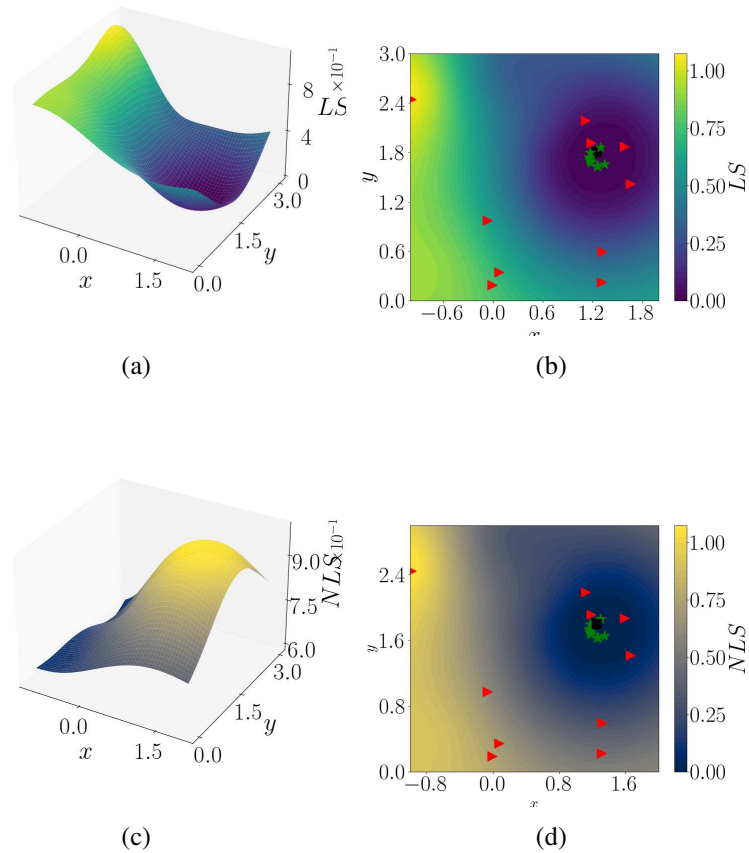


Figure 3.11: Parameter inference via Bayesian inversion for the Mixed Gaussian–Periodic benchmark function. Red triangles: initial guesses; black square: observed quantity of interest; green stars: MAP estimates. (a) LS surface. (b) LS contour. (c) NLS surface. (d) NLS contour.

both benchmark problems indicates that the adaptive procedure identifies such regions efficiently, without requiring a dense exploration of the entire parameter domain.

From the perspective of parameter inference, the two benchmarks exhibit complementary behaviours. The Mixed Gaussian–Periodic function gives rise to a non-convex inverse landscape with a dominant basin of attraction and a localized posterior distribution. Although the least-squares surface is not globally convex, the posterior remains sufficiently concentrated to support local identifiability and a meaningful characterization of uncertainty. The Rosenbrock benchmark, by contrast, represents a strongly anisotropic yet well-posed inverse problem, for which the surrogate-assisted Bayesian inversion procedure identifies a unique MAP estimate and a sharply concentrated posterior distribution. The resulting posterior geometry reflects the anisotropic structure of the forward response, indicating that the framework captures not only the location of the inferred solution but also the directional structure of the associated uncertainty.




Figure 3.12: Parameter inference via Bayesian inversion for the Rosenbrock benchmark function. Red triangles: initial guesses; black square: observed quantity of interest; green star: MAP estimate. (a) LS surface. (b) LS contour. (c) NLS surface. (d) NLS contour.

The two-dimensional results therefore complement the one-dimensional analysis in a fundamental way. Whereas the one-dimensional benchmarks primarily highlight multimodality, oscillatory behaviour, and localized nonlinearities, the two-dimensional setting shows that the framework remains reliable in the presence of parameter coupling and geometrically structured response surfaces. At the same time, the present analysis is intentionally restricted to low-dimensional benchmark problems and should not be interpreted as evidence of scalability to high-dimensional settings. Rather, it provides a controlled validation of the main methodological components of the proposed framework.

Overall, the results support three main conclusions. First, uncertainty-aware adaptive sampling enables the construction of accurate surrogate models with a limited number of high-fidelity evaluations. Second, the resulting surrogate models provide a reliable basis for Bayesian inversion and posterior analysis. Third, the framework enables a consistent characterization of

different inverse problem structures, including non-convex, multimodal, and anisotropic cases, through explicit posterior analysis while maintaining computational tractability in the considered benchmark settings.

3.1.4.2.3 Discussion

The numerical validation results presented in Sections 3.1.4.2.1 and 3.1.4.2.2 confirm the effectiveness of the proposed uncertainty-aware framework in two-dimensional inverse problem settings characterized by coupled parameters, non-convex response surfaces, and anisotropic features. The two-dimensional benchmarks extend the one-dimensional analysis (Section 3.1.4.1) by introducing non-trivial parameter interactions and geometrically structured response landscapes, while still providing a controlled setting in which the behaviour of the framework can be interpreted in a clear and unambiguous manner.

As in the one-dimensional case, the adaptive surrogate modeling strategy enables the construction of accurate Gaussian process surrogate models using a limited number of high-fidelity evaluations. The results obtained for both the Mixed Gaussian–Periodic and Rosenbrock benchmark functions show that uncertainty-aware sampling based on the predictive variance effectively concentrates evaluations in informative regions of the parameter space. This property becomes particularly relevant in two-dimensional settings, where the efficiency of uniform sampling deteriorates rapidly and the spatial distribution of the training points plays a central role in determining surrogate accuracy.

The numerical results further show that the quality of the surrogate model remains sufficient to support reliable Bayesian inversion, even in the presence of parameter coupling and non-trivial response geometries. In particular, the two benchmark functions highlight complementary inference behaviours. The Mixed Gaussian–Periodic function gives rise to a non-convex inverse landscape with a localized posterior structure, whereas the Rosenbrock function leads to a strongly anisotropic but well-defined posterior concentrated around a unique MAP estimate. These results indicate that the proposed framework is able to distinguish between different inverse problem structures and to provide an interpretable characterization of the associated epistemic uncertainty.

From a methodological perspective, the two-dimensional study reinforces the role of adaptive surrogate construction as a prerequisite for efficient inverse analysis. The quality of the inferred solution depends directly on the ability of the surrogate model to accurately reproduce the forward response in the regions that govern the least-squares functional and the posterior distribution. The observed convergence behaviour suggests that predictive-variance-based sampling provides an effective mechanism for identifying such regions without requiring dense sampling of the full parameter domain.

At the same time, the present results should be interpreted within the scope of the adopted benchmark setting. The analysis demonstrates the consistency and robustness of the proposed framework in low-dimensional problems, but it does not constitute evidence of scalability to higher-dimensional inverse problems. Rather, the two-dimensional benchmarks provide a controlled validation of the main methodological components of the framework and support their use in more structured application-oriented settings.

3.1.5 Conclusions

The present work introduces an uncertainty-aware framework for the solution of inverse problems, combining Gaussian process (GP) surrogate modeling with Bayesian inversion within a unified probabilistic setting. The proposed approach is designed for scenarios characterized by computationally expensive high-fidelity models and limited data availability, and provides a structured methodology for parameter inference together with a quantitative characterization of epistemic uncertainty.

The surrogate modeling stage relies on GP surrogate models that are iteratively refined through an adaptive sampling strategy driven by predictive uncertainty. This strategy enables an efficient exploration of the parameter space and allows the construction of accurate surrogate representations using a limited number of high-fidelity model evaluations. The resulting surrogate models preserve the essential features of the forward response and provide a computationally tractable approximation suitable for subsequent inverse analysis.

Once the surrogate model is constructed, the inverse problem is addressed within a Bayesian inversion framework. This formulation yields Maximum A Posteriori (MAP) estimates together with a local approximation of the posterior distribution of the input parameters. The probabilistic characterization of the solution enables the identification of admissible parameter regions, the assessment of parameter identifiability, and the quantification of epistemic uncertainty associated with the inverse problem.

The framework is validated on a set of one- and two-dimensional analytical benchmark functions exhibiting nonlinear, multimodal, and non-convex behaviour. These benchmarks provide a controlled environment to analyse the interplay between surrogate modeling accuracy and parameter inference. The numerical results indicate that, within the considered low-dimensional settings, the proposed approach provides accurate surrogate approximations, stable parameter estimates, and informative local characterizations of posterior uncertainty, while requiring a limited number of high-fidelity evaluations.

The results further highlight the role of surrogate accuracy in determining the quality of the inverse solution. In particular, the adaptive sampling strategy improves the representation of the forward model in regions characterized by high predictive uncertainty, which in turn contributes to a more reliable approximation of the posterior distribution. The framework is shown to capture both multimodal inverse problems, where multiple admissible solutions exist, and well-posed problems characterized by a unique and identifiable parameter configuration.

The present study is intentionally restricted to low-dimensional benchmark problems in order to provide a controlled and interpretable assessment of the methodological components of the framework. The extension to higher-dimensional parameter spaces introduces additional challenges related to sampling efficiency, surrogate model scalability, and posterior exploration, which are not addressed in the present work.

Future research will therefore focus on extending the framework to more complex inverse problems, including higher-dimensional parameter spaces, multi-fidelity modeling strategies, and physics-informed surrogate models. In addition, the development of scalable Bayesian inference techniques will be essential to maintain computational tractability in large-scale applications.

Overall, the proposed framework provides a computationally efficient and probabilistically consistent approach for surrogate-assisted Bayesian inversion, enabling parameter inference and uncertainty quantification in settings where direct evaluation of high-fidelity models is computationally prohibitive.

3.2 Uncertainty-Aware Bayesian Inversion Using Adaptive Gaussian-Process Surrogate Modeling for the Inverse Design of Spinodoid Architected Materials

3.2.1 Introduction

Inverse problems provide a mathematical framework for identifying unknown input parameters from observed model responses and arise in a wide range of engineering and scientific applications [36, 37, 128, 130, 178]. In the context of architected materials [10, 179, 180], such formulations naturally lead to inverse design problems, where the objective is to determine microstructural parameters that reproduce prescribed macroscopic properties [23, 181, 182]. These problems rely on a forward mapping that relates design parameters to effective material responses, typically evaluated through high-fidelity models based on computational homogenization, which capture the structure–property relationship of complex microstructures [25, 26, 183].

Despite their predictive capability, high-fidelity forward models remain computationally demanding [25, 184], particularly when repeatedly evaluated within inverse design procedures. This limitation becomes critical in probabilistic settings, where parameter inference requires extensive exploration of the admissible design space together with repeated evaluations of likelihood functions [185, 186]. To alleviate this computational burden, surrogate models are commonly employed to approximate the forward mapping. Among available approaches, Gaussian-process regression provides a probabilistic surrogate model representation that yields both predictive mean and predictive covariance, thereby quantifying epistemic uncertainty associated with limited training data [141, 145, 147, 170, 187, 188].

Within the uncertainty quantification community, several alternative surrogate modeling strategies have been developed [8, 141, 149]. Polynomial chaos expansions and sparse spectral representations provide efficient approximations for forward uncertainty propagation and sensitivity analysis [101], while general-purpose frameworks such as UQLab enable systematic surrogate-based modeling and probabilistic analysis in engineering applications [104, 189]. These developments establish surrogate modeling as a central component of modern uncertainty quantification methodologies.

The integration of surrogate models within inverse problems requires a coherent probabilistic formulation. Within a Bayesian inversion framework [157, 190], the posterior distribution of the unknown parameters is obtained by combining prior information with a likelihood function that measures the agreement between observed quantities and model predictions [36, 37, 157, 162, 191]. When a surrogate model replaces the high-fidelity forward model, the

likelihood depends on an approximate representation of the forward mapping. In this setting, two distinct sources of uncertainty influence the inference process: the predictive uncertainty associated with the surrogate model and the residual discrepancy between surrogate predictions and high-fidelity evaluations [6]. Neglecting either contribution may lead to biased parameter estimates and overconfident inference.

Discrepancy-aware Bayesian formulations introduced by Kennedy and O’Hagan [30] and further developed by Higdon et al. [192] provide a principled framework for incorporating model inadequacy within the likelihood. Subsequent developments in surrogate-based Bayesian inversion have investigated both spectral approximations [161] and Gaussian-process surrogate models for efficient parameter inference [193–195]. These contributions demonstrate that surrogate modeling, discrepancy modeling, and Bayesian inversion can be combined within a coherent probabilistic setting.

However, existing approaches often treat surrogate predictive uncertainty, residual model discrepancy, and systematic approximation bias as separate components, or incorporate them only partially within practical inversion workflows. As a result, their combined effect is not always represented explicitly within a single likelihood formulation suitable for complex engineering applications [27, 141, 196, 197]. In addition, practical implementations require adaptive surrogate construction strategies that remain computationally tractable while preserving consistency with the underlying high-fidelity model [154, 155].

In the context of architected materials, and in particular for spinodoid architected materials [21, 23, 198, 199], the inverse design problem is further complicated by the nonlinear and anisotropic nature of the structure–property relationship. The effective mechanical response depends on geometric parameters controlling anisotropy and volume fraction and is evaluated through computational homogenization [26, 183, 200]. The resulting mapping between design parameters and effective properties is highly nonlinear and may exhibit strong parameter interactions, thereby increasing the sensitivity of the inverse problem to approximation errors in the forward model.

The present work develops an uncertainty-aware Bayesian inversion framework for the inverse design of spinodoid architected materials, with emphasis on the construction of a likelihood model that explicitly accounts for surrogate-induced uncertainty. A Gaussian-process surrogate model is constructed through an adaptive sampling strategy driven by predictive uncertainty, with the objective of efficiently approximating the structure–property mapping over the admissible design space. The surrogate model is embedded within the inversion procedure through a likelihood formulation that incorporates both predictive mean and predictive covariance.

Residual discrepancy between surrogate predictions and high-fidelity evaluations is represented through an effective covariance matrix estimated from validation residuals, providing a data-driven approximation of the residual modeling error. In addition, a bias-corrected surrogate model is introduced to reduce systematic discrepancies in the predictive mean. The resulting formulation combines surrogate predictive uncertainty and residual discrepancy within a single likelihood model tailored to surrogate-based inversion.

The methodology is assessed on synthetic inverse design problems defined over reduced two- and four-dimensional parameter subspaces, where target responses are generated from

the high-fidelity forward model. Parameter inference is performed via Maximum A Posteriori estimation. This controlled setting enables the individual contributions of predictive uncertainty, discrepancy modeling, and bias correction to be analyzed in isolation, and is intended to assess the internal consistency and numerical behavior of the proposed formulation rather than to reproduce the full complexity of experimental inverse-design scenarios.

The manuscript is organized as follows. Section 3.2.2 describes the construction and parametrization of spinodoid architected materials. Section 3.2.3 presents the high-fidelity forward model and the construction of the structure–property dataset. Section 3.2.4 details the uncertainty-aware Bayesian inversion framework, including adaptive surrogate construction and likelihood formulation. Section 3.2.5 reports the results of the inverse design analyses in reduced parameter spaces. Finally, Section 3.2.5.3 summarizes the main findings and outlines directions for future work.

3.2.2 Spinodoid Architected Materials

Spinodoid architected materials define a class of bicontinuous three-dimensional microstructures characterized by smooth solid–void interfaces, high geometric continuity, and strong topological connectivity [21–23, 198, 199]. Such features induce complex structure–property relationships and make spinodoid architectures particularly suitable for inverse design formulations.

The term *spinodoid* refers to morphologies inspired by the spatial configurations arising during the early stages of spinodal decomposition in binary mixtures [21, 201–204]. Within the Cahn–Hilliard theory [201, 205, 206], small fluctuations in concentration evolve into interconnected bicontinuous patterns characterized by smooth interfaces and statistical homogeneity. Direct numerical simulation of the governing evolution equations is computationally demanding and therefore unsuitable for design-oriented settings requiring repeated evaluations.

To enable efficient parametrization, spinodoid architected materials are generated through stochastic representations based on stationary Gaussian random fields [21, 23, 207]. Within this framework, the morphology is defined through a scalar field constructed as a superposition of harmonic wave modes with identical wavelength and randomly distributed orientations and phase shifts.

Let $\Omega \subset \mathbb{R}^3$ denote a cubic reference domain with spatial coordinate $\mathbf{x} = (X, Y, Z)^\top$. The scalar field is represented as

$$\varphi(\mathbf{x}) = \sqrt{\frac{2}{n_{\text{wave}}}} \sum_{k=1}^{n_{\text{wave}}} \cos(\mathbf{k}_k \cdot \mathbf{x} + \phi_k), \quad (3.19)$$

where n_{wave} denotes the number of wave modes, $\mathbf{k}_k \in \mathbb{R}^3$ are wave vectors of identical magnitude, and $\phi_k \in [0, 2\pi)$ are uniformly distributed phase shifts. The phase angles control spatial translations, while the orientations of the wave vectors govern directional correlations. For sufficiently large n_{wave} , the resulting field converges (in distribution) to a stationary Gaussian random field exhibiting smooth and statistically homogeneous fluctuations consistent with spinodal-like morphologies.

Directional anisotropy is introduced by restricting the admissible orientations of the wave vectors. Let \mathbf{e}_i , $i = 1, 2, 3$, denote the canonical basis of \mathbb{R}^3 and let $\widehat{\mathbf{k}} = \mathbf{k}/\|\mathbf{k}\|$ denote the normalized direction. Three anisotropy angles $\theta_i \in [0, \pi/2]$ define conical regions around each axis such that

$$\left| \widehat{\mathbf{k}} \cdot \mathbf{e}_i \right| \geq \cos(\theta_i). \quad (3.20)$$

Small values of θ_i enforce alignment along the corresponding axis and produce strongly anisotropic morphologies, whereas increasing θ_i progressively relaxes the constraint and leads to more isotropic configurations. The limiting case $\theta_i = \pi/2$ corresponds to the absence of directional constraints.

The scalar field $\varphi(\mathbf{x})$ is converted into a binary solid–void architecture through level-set thresholding [21, 23]. Given a threshold $\tau \in \mathbb{R}$, the solid and void domains are defined as

$$\Omega_s(\tau) = \{\mathbf{x} \in \Omega : \varphi(\mathbf{x}) \geq \tau\}, \quad \Omega_v(\tau) = \Omega \setminus \Omega_s(\tau), \quad (3.21)$$

with associated volume fraction

$$v_f(\tau) = \frac{|\Omega_s(\tau)|}{|\Omega|}. \quad (3.22)$$

For a prescribed $v_f \in (0, 1)$, the threshold τ is selected such that $v_f(\tau) = v_f$, thereby enabling independent control of density and anisotropy. The resulting interface corresponds to an iso-surface of φ , yielding geometrically smooth and continuous structures.

Global orientation is controlled through rigid rotations. Let $\mathbf{R}(\psi_1, \psi_2, \psi_3) \in \text{SO}(3)$ denote the rotation matrix associated with three Euler angles. The wave vectors are transformed as

$$\mathbf{k}_k \leftarrow \mathbf{R}(\psi_1, \psi_2, \psi_3)\mathbf{k}_k, \quad k = 1, \dots, n_{\text{wave}}, \quad (3.23)$$

which induces a rigid rotation of the generated microstructure without altering its topology or characteristic length scales.

The complete generation process is illustrated in Figure 3.13, which summarizes the construction of the scalar field, the introduction of anisotropy, the level-set thresholding, and the global orientation control.

Within the present parametrization, each spinodoid configuration is represented by the parameter vector

$$\Theta = (\theta_1, \theta_2, \theta_3, v_f, \psi_1, \psi_2, \psi_3)^\top \in \mathcal{S}_\Theta \subset \mathbb{R}^7, \quad (3.24)$$

where the admissible design space is defined by $\theta_i \in [0, \pi/2]$, $v_f \in (0, 1)$, and $\psi_i \in [0, 2\pi)$.

Different regions of the parameter space correspond to distinct canonical morphologies. Lamellar, columnar, cubic, and isotropic configurations emerge as limiting cases associated with specific combinations of anisotropy angles and exhibit markedly different connectivity patterns, which directly influence the effective mechanical response.

The parametrization introduced in this section defines a mapping from the design space \mathcal{S}_Θ to the corresponding microstructural geometry. This mapping provides the geometric input to the high-fidelity forward model described in the following section, where the associated effective properties are evaluated.

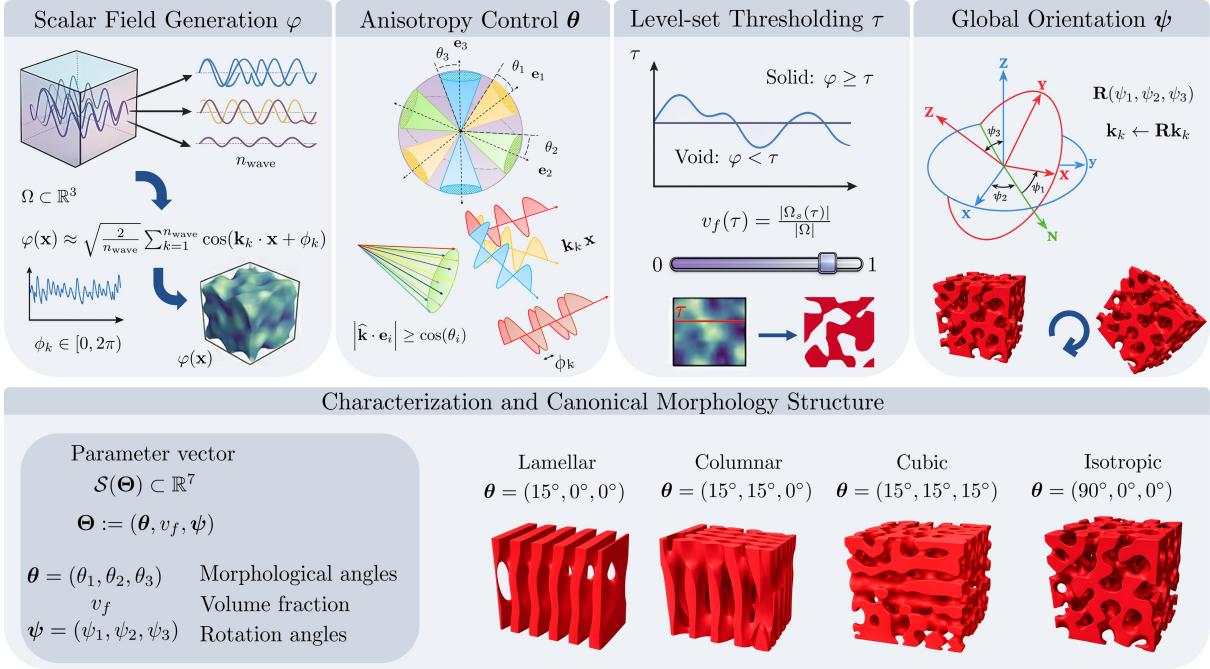


Figure 3.13: Generation and parametrization of spinodoid architected materials. From left to right, the upper part illustrates the construction process: generation of the scalar Gaussian random field; restriction of wave-vector orientations via anisotropy angles θ_i ; level-set thresholding defining the solid domain and controlling the volume fraction v_f ; and rigid rotation τ governed by the angles ψ_i . The lower part presents representative morphologies associated with characteristic anisotropy configurations, including lamellar, columnar, cubic, and isotropic structures. The complete microstructural description is defined by the parameter vector $\Theta = (\theta_1, \theta_2, \theta_3, v_f, \psi_1, \psi_2, \psi_3)$.

3.2.3 High-Fidelity Forward Model

The parametrization introduced in Section 3.2.2 defines a mapping from the design space \mathcal{S}_Θ to the corresponding class of voxelized spinodoid architected materials. The evaluation of the associated effective mechanical response requires the solution of a multiscale boundary value problem, which is addressed through computational homogenization [26, 183, 200].

Within the present framework, the high-fidelity forward model defines a deterministic mapping between the morphological parameters and the corresponding effective elastic response,

$$\mathcal{F} : \mathcal{S}_\Theta \rightarrow \mathbb{R}^3, \quad \Theta \mapsto \mathbf{y}_{\text{HF}}, \quad (3.25)$$

where \mathbf{y}_{HF} denotes the vector of effective stiffness components defined in Equation (3.28). The mapping \mathcal{F} constitutes the reference forward model employed in the uncertainty-aware Bayesian inversion framework introduced in Section 3.2.4. The deterministic character of \mathcal{F} is understood here in the sense of the modeling convention adopted in the present work, where a single microstructural realization is retained for each admissible parameter vector, as detailed in Section 3.2.3.1.

All high-fidelity simulations are performed using the Fourier-Accelerated Nodal Solver (FANS) [208–211], an FFT-based computational framework specifically designed for voxelized heterogeneous materials. FFT-based homogenization methods exploit the regular grid structure of discretized microstructures and provide efficient solution strategies for periodic boundary value problems [209, 210, 212]. Within the present work, FANS defines the numerical realization of the high-fidelity forward model \mathcal{F} .

For each admissible parameter vector $\Theta \in \mathcal{S}_\Theta$ (Equation (3.24)), the numerical workflow consists of four steps: (i) generation of the corresponding spinodoid architected material, (ii) numerical homogenization of the representative volume element, (iii) extraction of selected effective mechanical properties, and (iv) construction of the structure–property dataset. Each component is described in the following.

3.2.3.1 Generation of Spinodoid Geometry

For each parameter configuration Θ , the corresponding spinodoid architected material is generated according to the Gaussian-random-field construction described in Section 3.2.2 [21, 23]. The microstructure is defined by the scalar field $\varphi(\mathbf{x})$ introduced in Equation (3.19), whose directional correlations are controlled by the anisotropy angles θ_i through the conical constraints in Equation (3.20).

The binary solid–void architecture is obtained through the level-set thresholding procedure in Equation (3.21), with the threshold selected such that the prescribed solid volume fraction v_f is satisfied according to Equation (3.22). Global orientation is introduced through the rigid rotation defined in Equation (3.23).

The scalar field is evaluated on a regular Cartesian grid discretizing the domain Ω . In the present work, the microstructure is represented on a voxel grid of resolution 128^3 , which provides an accurate representation of the smooth interfaces characteristic of spinodoid architected materials while maintaining computational tractability.

The construction procedure involves stochastic sampling of wave-vector orientations and phase shifts. Consequently, the mapping between the design parameters and the corresponding effective response is, in principle, stochastic. In the present work, however, a single realization is retained for each parameter configuration in order to define an operationally deterministic mapping between Θ and the effective response, consistent with the definition of the forward model \mathcal{F} . This choice enables the construction of a well-defined surrogate model and the formulation of the inverse problem in a tractable setting.

A systematic quantification of the stochastic variability associated with microstructure generation and its impact on the effective properties is not considered in the present study and is left for future investigation.

The resulting voxelized geometries are stored in HDF5 format and subsequently used as input for the homogenization analysis.

3.2.3.2 Numerical Homogenization

The effective elastic response of each generated microstructure is computed through numerical homogenization performed on the corresponding representative volume element. The analysis is conducted under the assumptions of infinitesimal strains and linear elasticity.

The solid phase is modeled as an isotropic linear elastic material characterized by Young's modulus E and Poisson's ratio ν . The pore space is represented as a highly compliant phase with stiffness several orders of magnitude smaller than that of the solid phase, enabling its numerical treatment within the FFT framework [209, 210, 213].

The elastic properties of the solid phase are expressed in terms of the bulk modulus K and shear modulus G , given by

$$K = \frac{E}{3(1-2\nu)}, \quad G = \frac{E}{2(1+\nu)}. \quad (3.26)$$

Normalized material properties are adopted by setting $E = 1$ and $\nu = 0.3$, yielding dimensionless effective stiffness components while preserving relative comparisons among different microstructures.

Periodic boundary conditions are imposed on the representative volume element. Macroscopic loading is introduced through prescribed uniform macroscopic strain tensors, and the corresponding equilibrium problems are solved iteratively using a conjugate-gradient scheme with an absolute residual tolerance of 10^{-10} and a maximum of 100 iterations per load case.

The homogenization procedure yields the volume-averaged macroscopic stress response associated with prescribed strain states, from which the effective elasticity tensor \mathbf{C}^{eff} is determined through the constitutive relation

$$\boldsymbol{\sigma}^{\text{M}} = \mathbf{C}^{\text{eff}} : \boldsymbol{\varepsilon}^{\text{M}}. \quad (3.27)$$

3.2.3.3 Extraction of Effective Mechanical Properties

The inverse design formulation focuses on selected components of the effective elasticity tensor associated with normal deformation along the principal Cartesian directions. The quantities of interest are defined as

$$\mathbf{y}_{\text{HF}} = \left(C_{1111}^{\text{eff}}, C_{2222}^{\text{eff}}, C_{3333}^{\text{eff}} \right)^{\text{T}}. \quad (3.28)$$

These components provide a compact and physically meaningful representation of the anisotropic elastic response induced by the morphological parameters $\boldsymbol{\Theta}$. Variations in the anisotropy angles θ_i , the volume fraction v_f , and the global orientation ψ_i modify the directional connectivity of the solid phase and therefore affect the stiffness response along the principal directions.

The restriction to the components in Equation (3.28) defines a reduced output space that remains sensitive to the dominant anisotropic features of the microstructure while preserving a tractable inverse design formulation.

3.2.3.4 Construction of the Structure–Property Dataset

Repeated evaluations of the high-fidelity forward model over the design space generate a dataset relating morphological parameters to effective mechanical properties. For each configuration $\Theta_i \in \mathcal{S}_\Theta$, the homogenization analysis produces an output vector $\mathbf{y}_i \in \mathbb{R}^3$, such that

$$\mathcal{D} = \{(\Theta_i, \mathbf{y}_i)\}_{i=1}^N, \quad \mathbf{y}_i = \mathcal{F}(\Theta_i). \quad (3.29)$$

The dataset \mathcal{D} provides the high-fidelity evaluations required for the construction of Gaussian-process surrogate models and defines the reference information employed within the likelihood formulation of the Bayesian inversion framework described in Section 3.2.4.

The numerical framework described in this section establishes the mapping between the design parameters and the corresponding effective properties and therefore provides the physical foundation of the inverse design problem addressed in the present work.

3.2.4 Uncertainty-Aware Bayesian Inversion Framework

The high-fidelity forward model introduced in Section 3.2.3 defines the structure–property mapping in Equation (3.25), while the corresponding dataset of high-fidelity evaluations is given in Equation (3.29). The inverse design problem considered in the present work consists in identifying parameter configurations $\Theta \in \mathcal{S}_\Theta$ that reproduce a prescribed target response in the space of effective properties defined in Equation (3.28).

Direct solution of this problem through repeated evaluations of the high-fidelity forward model is computationally demanding, as each evaluation requires the numerical homogenization procedure described in Section 3.2.3. This limitation motivates the introduction of a surrogate-model-based inversion framework, in which the forward mapping is approximated while preserving a coherent probabilistic formulation of the inverse problem.

The proposed framework combines adaptive Gaussian-process surrogate modeling with Bayesian inversion. Gaussian-process surrogate models provide a probabilistic approximation of the forward mapping, while Bayesian inversion enables parameter inference by combining prior information with a likelihood-based data misfit [36, 37, 214]. Particular attention is devoted to the consistent and explicit treatment of uncertainty introduced by the surrogate approximation. The formulation explicitly accounts for surrogate predictive uncertainty and residual discrepancy with respect to the high-fidelity forward model, following the principles of discrepancy-aware Bayesian calibration [30, 141, 192, 215, 216].

The overall workflow follows a two-stage structure. A Gaussian-process surrogate model is first constructed through adaptive sampling driven by predictive uncertainty. The resulting probabilistic approximation of the forward mapping is then embedded within a Bayesian inversion procedure, where the likelihood formulation incorporates both surrogate predictive covariance and a data-driven estimate of residual discrepancy. This construction enables the use of surrogate models within the inversion process while maintaining consistency with the underlying high-fidelity forward model.

3.2.4.1 Adaptive Construction of Gaussian-Process Surrogate Models

The surrogate modeling stage starts from the high-fidelity dataset \mathcal{D} defined in Equation (3.29). The objective is to construct a probabilistic approximation of the forward mapping over the admissible design space using a limited number of high-fidelity evaluations.

Gaussian-process regression defines a distribution over functions and yields, at each point in the parameter space, both a predictive mean and a predictive covariance [29, 146, 217]. Since the forward model output is vector-valued (see Equation (3.28)), an independent multi-output construction is adopted, where one Gaussian-process surrogate model is trained for each component. This choice provides numerical robustness, simplifies hyperparameter estimation, and yields direct access to output-wise predictive variances required for the likelihood formulation. Cross-output dependence is not modeled explicitly at the surrogate level in order to avoid the additional complexity associated with structured multi-output kernels, which would significantly increase computational cost and complicate hyperparameter estimation. Instead, residual cross-output dependencies are incorporated at the inversion stage through the effective covariance matrix Σ_{eff} .

For $k = 1, 2, 3$, the k -th response component is modeled as

$$\hat{y}_k(\Theta) \sim \mathcal{GP}(m_k(\Theta), k_k(\Theta, \Theta')), \quad (3.30)$$

which yields the posterior predictive distribution

$$\hat{y}_k(\Theta) \mid \mathcal{D} \sim \mathcal{N}\left(\mu_k(\Theta), \sigma_k^2(\Theta)\right). \quad (3.31)$$

The predictive mean and standard deviation are collected as

$$\mu(\Theta) = (\mu_1(\Theta), \mu_2(\Theta), \mu_3(\Theta))^\top, \quad (3.32)$$

$$\sigma(\Theta) = (\sigma_1(\Theta), \sigma_2(\Theta), \sigma_3(\Theta))^\top. \quad (3.33)$$

Input variables are mapped to the unit hypercube, and outputs are standardized independently to improve numerical conditioning. Covariance kernels are selected from standard families, such as Matérn kernels with automatic relevance determination (ARD), allowing different characteristic length scales along each input dimension.

The surrogate model is constructed adaptively. Starting from an initial training set, additional samples are selected iteratively by maximizing an acquisition score over a candidate set

$$\mathcal{X}_{\text{cand}} = \left\{ \Theta^{(j)} \right\}_{j=1}^{N_{\text{cand}}}. \quad (3.34)$$

The uncertainty measure is defined as

$$U_{\text{IVAR}}(\Theta) = \sum_{k=1}^3 \sigma_k^2(\Theta), \quad (3.35)$$

and the adaptive selection criterion is given by

$$S(\Theta) = \tilde{U}_{\text{IVAR}}(\Theta) \tilde{d}(\Theta)^{\beta_d}. \quad (3.36)$$

At each iteration, the selected candidate is evaluated using the high-fidelity forward model, appended to the dataset, and used to retrain the surrogate model. Predictive performance is assessed on independent validation and test sets through standard regression metrics.

3.2.4.2 Formulation of the Bayesian Inversion Problem

Let $\mathbf{y}^* \in \mathbb{R}^3$ denote the target response, corresponding to the vector of effective mechanical properties defined in Equation (3.28). In the present setting, the target is generated from the high-fidelity forward model for a reference parameter configuration, thereby defining a controlled inverse problem in which the ground-truth parameters are known.

Within the Bayesian framework, the posterior density of the unknown parameter vector $\Theta \in \mathcal{S}_\Theta$ is defined as

$$\pi(\Theta | \mathbf{y}^*) \propto \pi(\mathbf{y}^* | \Theta) \pi(\Theta), \quad (3.37)$$

where $\pi(\mathbf{y}^* | \Theta)$ denotes the likelihood function and $\pi(\Theta)$ represents the prior distribution. In the present work, a uniform prior is adopted over the admissible design space \mathcal{S}_Θ , reflecting the absence of additional prior information and enforcing only bound constraints on the parameters.

The likelihood explicitly accounts for both surrogate predictive uncertainty and residual approximation error with respect to the high-fidelity forward model. The conditional model is defined as

$$\mathbf{y}^* | \Theta \sim \mathcal{N}(\mu(\Theta), \Sigma_{\text{tot}}(\Theta)), \quad (3.38)$$

where $\mu(\Theta) \in \mathbb{R}^3$ denotes the predictive mean of the Gaussian-process surrogate model evaluated at Θ , and $\Sigma_{\text{tot}}(\Theta) \in \mathbb{R}^{3 \times 3}$ represents the total predictive covariance.

The total covariance is defined as

$$\Sigma_{\text{tot}}(\Theta) = \Sigma_{\text{eff}} + \Sigma_{\text{GP}}(\Theta), \quad (3.39)$$

where the first term $\Sigma_{\text{eff}} \in \mathbb{R}^{3 \times 3}$ is an empirical covariance matrix, while the second term $\Sigma_{\text{GP}}(\Theta)$ represents the predictive covariance provided by the surrogate model. In particular,

$$\Sigma_{\text{GP}}(\Theta) = \text{diag}\left(\sigma_1^2(\Theta), \sigma_2^2(\Theta), \sigma_3^2(\Theta)\right), \quad (3.40)$$

where $\sigma_k^2(\Theta)$ denotes the predictive variance of the k -th Gaussian-process surrogate model. This term captures the epistemic uncertainty associated with the surrogate approximation due to limited training data.

The matrix Σ_{eff} is estimated from out-of-sample residuals between high-fidelity evaluations and surrogate predictions on an independent validation dataset. In the present formulation, Σ_{eff} is assumed independent of Θ , thereby providing a global approximation of residual model discrepancy over the admissible parameter space, accounting for both systematic bias and

residual variability not captured by the surrogate model. This assumption provides a tractable approximation and is consistent with standard discrepancy-aware Bayesian formulations, in which a global error model is adopted in the absence of sufficient data to resolve parameter-dependent discrepancies. This term introduces cross-output correlations within the likelihood model, so that the total covariance $\Sigma_{\text{tot}}(\Theta)$ provides a unified representation of both local predictive uncertainty and global model discrepancy. The additive structure of $\Sigma_{\text{tot}}(\Theta)$ assumes statistical independence between surrogate predictive uncertainty and residual model discrepancy, while cross-output correlations are introduced through Σ_{eff} .

Under these assumptions, the negative log-likelihood is given by

$$\mathcal{L}(\Theta) = \frac{1}{2}(\mathbf{y}^* - \mu(\Theta))^\top \Sigma_{\text{tot}}(\Theta)^{-1}(\mathbf{y}^* - \mu(\Theta)) + \frac{1}{2} \log \det \Sigma_{\text{tot}}(\Theta) + \frac{3}{2} \log(2\pi), \quad (3.41)$$

which corresponds to the standard expression of the negative log-density of a multivariate Gaussian distribution in dimension three.

To mitigate systematic approximation errors in the predictive mean, an optional bias-corrected surrogate model can be introduced as

$$\mu_{\text{BC}}(\Theta) = \mu(\Theta) + \mu_\delta(\Theta), \quad (3.42)$$

where $\mu_\delta(\Theta)$ denotes a data-driven correction term constructed from residual data using a secondary surrogate model trained on out-of-sample residuals. This correction aims at reducing systematic bias in the surrogate prediction while preserving its probabilistic structure. The correction term $\mu_\delta(\Theta)$ affects only the predictive mean, while residual uncertainty remains captured by Σ_{eff} , thereby avoiding double counting of uncertainty contributions within the likelihood model.

Parameter inference is performed via Maximum A Posteriori (MAP) estimation,

$$\Theta_{\text{MAP}} = \arg \min_{\Theta \in \mathcal{S}_\Theta} [\mathcal{L}(\Theta) - \log \pi(\Theta)]. \quad (3.43)$$

Under a uniform prior over \mathcal{S}_Θ , the MAP estimate reduces to the minimization of the negative log-likelihood over the admissible domain.

The resulting optimization problem is generally non-convex due to the nonlinear dependence of the surrogate model on the parameters. To improve robustness with respect to local minima, a multi-start strategy is employed. In addition, a posterior-focused refinement stage can be introduced by selecting new samples according to

$$S_{\text{AR}}(\Theta) \propto \exp\left(-\frac{\Phi(\Theta) - \Phi_{\min}}{T}\right) U_{\text{IVAR}}(\Theta) \tilde{d}(\Theta)^{\beta_{\text{AR}}}, \quad (3.44)$$

where $\Phi(\Theta) = \mathcal{L}(\Theta)$ denotes the negative log-likelihood, Φ_{\min} denotes the minimum value of Φ over the explored parameter set, $U_{\text{IVAR}}(\Theta)$ is the integrated predictive variance defined in Equation (3.35), $\tilde{d}(\Theta)$ is a normalized distance from the current training set, β_{AR} controls the

balance between exploration and exploitation, and T is a temperature parameter regulating the concentration around high-probability regions.

The formulation remains consistent with the high-fidelity forward model, as the surrogate model is constructed exclusively from high-fidelity evaluations and the likelihood explicitly accounts for surrogate-induced uncertainty and residual discrepancy within a unified probabilistic framework.

3.2.5 Results and Discussion

The present section assesses the performance of the uncertainty-aware Bayesian inversion framework introduced in Section 3.2.4 through controlled numerical experiments. The analysis focuses on the inverse design of spinodoid architected materials described in Section 3.2.2, with the objective of evaluating surrogate model accuracy, inversion consistency, and the impact of the uncertainty-aware likelihood formulation on parameter inference.

The inverse design problems are formulated over reduced parameter subspaces involving the anisotropy angles θ_i and the volume fraction v_f . These parameters govern the morphology of spinodoid architected materials and directly influence their effective mechanical response [21–23, 198, 199]. The restriction to reduced parameter spaces enables a controlled assessment of the proposed framework while preserving the essential nonlinear structure–property mapping.

The rotation angles ψ_i are not included in the inversion parameter space. As discussed in Section 3.2.2, these parameters define rigid-body rotations of the generated microstructure. Within the present framework, the effective response is interpreted with respect to intrinsic material directions, so that variations in ψ_i correspond to equivalent representations of the same structure–property mapping. Under this interpretation, the rotation parameters do not introduce additional independent variability in the forward operator and therefore do not contribute to the inverse identification problem.

The dataset is constructed according to the high-fidelity modeling procedure described in Section 3.2.3. For each parameter configuration, a single realization of the stochastic microstructure generation process is retained, thereby defining an operationally deterministic forward operator \mathcal{F} within the present study. This modeling choice ensures consistency with the surrogate-based inversion framework introduced in Section 3.2.4, in which the forward mapping is approximated by a deterministic surrogate model. The analysis therefore isolates the effects of surrogate approximation error and likelihood modeling, without introducing additional variability associated with stochastic microstructure generation. A systematic quantification of such stochastic variability is not considered in the present study and is left for future work.

Two configurations are investigated. The first corresponds to a two-dimensional parameter space defined by a subset of anisotropy parameters. The second considers a four-dimensional parameter space including the three anisotropy angles together with the volume fraction. These settings enable the proposed framework to be assessed under increasing levels of parametric complexity while maintaining a controlled inversion setting. This structure enables the respective roles of surrogate accuracy, predictive uncertainty, and discrepancy-aware likelihood modeling to be analyzed in a systematic and interpretable manner.

3.2.5.1 Results in Two-Dimensional Parameter Space

The first set of numerical experiments considers an inverse design problem defined over a two-dimensional parameter space, where the design vector is given by the anisotropy angles $\Theta = (\theta_1, \theta_2)$. The remaining parameters are fixed to reference values, namely $\theta_3 = 45^\circ$ and $\nu_f = 0.5$, consistently with the parametrization introduced in Section 3.2.2. The admissible domain is defined by $\theta_1, \theta_2 \in [0^\circ, 90^\circ]$.

The inverse problem is formulated in a synthetic setting by selecting a reference configuration $\Theta^* = (27^\circ, 34^\circ)$ and defining the target response \mathbf{y}^* through evaluation of the high-fidelity forward model \mathcal{F} (Equation (3.25)). This setting enables the consistency of the inversion procedure to be assessed in a controlled environment, where the ground-truth parameters are known and the behavior of the likelihood formulation can be analyzed. No observational noise is introduced and no additional physical model discrepancy is prescribed beyond that induced by the surrogate model approximation. Consequently, the discrepancy term included in the likelihood formulation reflects only residual errors associated with the surrogate approximation. More general scenarios involving observational uncertainty and model mismatch are beyond the scope of the present study and are left for future investigation.

The surrogate model is constructed using the adaptive Gaussian-process procedure described in Section 3.2.4.1. The initial training set consists of 8 samples, and the dataset is enriched over 60 adaptive iterations using a selection criterion that combines predictive variance and distance in the normalized parameter space. Independent Gaussian-process surrogate models are trained for each output component using a Matérn kernel with smoothness parameter $\nu = 2.5$. Predictive performance is assessed on independent validation and test sets of 60 samples each, generated through independent Latin hypercube sampling of the admissible domain.

Parameter inference is performed in a MAP setting as described in Section 3.2.4.2. A multistart strategy with 15 initial points is employed, combining basin-hopping with local L-BFGS-B minimization. The likelihood incorporates both surrogate predictive uncertainty and an effective covariance matrix Σ_{eff} (Equation (3.39)), estimated from an independent calibration dataset not used for surrogate model training, thereby reducing potential bias in the estimation of residual discrepancies. This construction enables the total covariance Σ_{tot} to capture both local predictive uncertainty and residual discrepancies with respect to high-fidelity evaluations.

To assess the impact of discrepancy estimation, the effective covariance matrix is computed using calibration datasets of increasing size (20, 50, 100, and 200 samples). This analysis enables the sensitivity of the inversion results to the accuracy of discrepancy modeling to be evaluated and provides insight into the trade-off between calibration cost and the robustness of the inferred parameters.

A uniform prior distribution is adopted over the admissible domain. Posterior-focused active refinement is performed over 5 iterations using the criterion defined in Equation (3.44), with the objective of improving surrogate accuracy in regions of high posterior probability. In addition, a bias-corrected surrogate model is constructed using a secondary Gaussian process trained on cross-validated residuals (Equation (3.42)), enabling the effect of systematic error correction in the predictive mean to be assessed within the inversion procedure. This construction

enables a direct comparison between standard and bias-corrected formulations within a consistent probabilistic framework.

3.2.5.1.1 Adaptive Construction of Gaussian-Process Surrogate Models

The performance of the Gaussian-process surrogate model in the two-dimensional parameter space is evaluated in terms of convergence behavior, predictive accuracy, and consistency with the underlying high-fidelity mapping.

The surrogate model approximates the forward relation defined in Equation (3.28), with output vector $\mathbf{y} = (C_{1111}^{\text{eff}}, C_{2222}^{\text{eff}}, C_{3333}^{\text{eff}})$. Given the independent multi-output construction adopted in Section 3.2.4.1, the analysis is carried out both at the level of individual response components and in terms of aggregated performance metrics.

The convergence of the adaptive construction is assessed component-wise through validation error metrics. Figure 3.14 reports the evolution of the normalized root mean square error (NRMSE) for each output component as a function of the number of training samples. All components exhibit a similar convergence pattern characterized by three distinct regimes. In the initial phase, corresponding to the first enrichment iterations, relatively large errors are observed due to the limited coverage of the parameter space. This regime is followed by a rapid error reduction occurring within approximately 10–20 samples, where the adaptive sampling strategy, driven by the integrated predictive variance defined in Equation (3.35), effectively identifies high-uncertainty regions and improves the global approximation. Beyond this transition, the convergence rate decreases and a stable regime is reached, with NRMSE values approaching 10^{-3} – 10^{-4} .

Despite the overall consistency of the convergence behavior, moderate differences are observed across the three components during the early stages of the adaptive process. In particular, the initial error levels and transient oscillations differ slightly, reflecting variations in the sensitivity of each effective modulus to the underlying design parameters. These discrepancies progressively vanish as the sampling density increases, indicating that the adaptive procedure successfully resolves anisotropic features of the structure–property mapping.

A consistent convergence behavior is observed for the coefficient of determination, as shown in Figure 3.15. All components exhibit a rapid increase toward unity, with R^2 values exceeding 0.99 after approximately 10–15 samples and stabilizing close to $R^2 \approx 1$. This behavior confirms that the surrogate model accurately captures the functional dependence of each output component over the admissible domain.

The global performance of the surrogate model is summarized in Figure 3.16, where aggregated NRMSE and R^2 metrics are reported. The combined error exhibits a rapid decrease of more than one order of magnitude during the early enrichment phase, followed by a gradual convergence toward a final value of $\text{NRMSE} \approx 4.6 \times 10^{-4}$. Similarly, the combined R^2 stabilizes at $R^2 \approx 0.999996$, indicating an almost exact reconstruction of the high-fidelity mapping within the resolution of the validation dataset.

The close agreement between component-wise and aggregated metrics indicates that the surrogate model provides a uniformly accurate approximation across all outputs. In addition, the absence of divergence between validation and test performance (not shown for brevity) suggests

that overfitting effects remain negligible and that the surrogate model generalizes reliably across the parameter space.




Figure 3.14: Validation NRMSE for each output component during adaptive enrichment. The error exhibits an initial transient regime followed by rapid decay and eventual stabilization, indicating progressive improvement of the surrogate approximation across all components.




Figure 3.15: Validation R^2 for each output component. All components rapidly approach unity, confirming the capability of the surrogate model to reproduce the high-fidelity mapping with high accuracy.

The surrogate model reconstruction over the (θ_1, θ_2) domain is shown in Figure 3.17. The model captures the heterogeneous behavior of the effective stiffness components and preserves the main structural features of the high-fidelity response.

The component C_{1111}^{eff} exhibits a strongly anisotropic behavior, with pronounced variations along θ_1 and a smoother dependence on θ_2 . The component C_{2222}^{eff} displays a more complex topology characterized by a saddle-like structure in the parameter space. In contrast, C_{3333}^{eff} shows a predominantly monotonic trend with reduced nonlinearity.

Across all output components, no spurious oscillations or numerical artifacts are observed, indicating that the surrogate model provides a stable and numerically consistent approximation of the forward mapping.

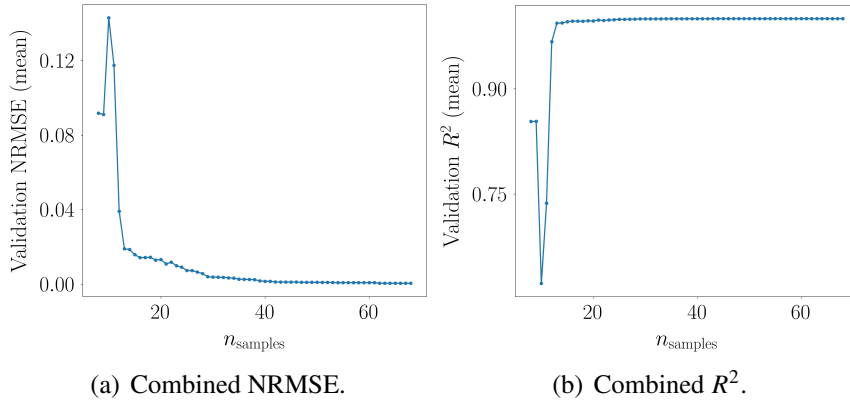


Figure 3.16: Aggregated validation metrics for the multi-output surrogate model. The combined indicators confirm the global accuracy and stability of the adaptive Gaussian-process approximation over the entire parameter space.




Figure 3.17: Gaussian-process surrogate model reconstruction of the effective stiffness components over the (θ_1, θ_2) domain.

Overall, the surrogate model achieves high predictive accuracy and provides a stable approximation of the high-fidelity mapping. Such a level of accuracy is essential for the subsequent inversion stage, where both the predictive mean and the predictive variance directly enter the likelihood formulation through Σ_{tot} and therefore directly influence the inferred posterior distribution via the likelihood.

3.2.5.1.2 Uncertainty-Aware Bayesian Inversion Results

The performance of the uncertainty-aware Bayesian inversion framework is assessed by analyzing the geometry of the objective functions, the role of the uncertainty model, and the accuracy and stability of the inferred solution.

The target response is defined in the space of effective properties introduced in Equation (3.28) and is obtained by evaluating the high-fidelity forward model at a prescribed reference configuration Θ^* . The inverse problem consists in identifying the parameter vector $\Theta = (\theta_1, \theta_2)$ through minimization of the MAP objective defined in Equation (3.43).

The geometry of the inverse problem is first examined through contour representations of the least-squares (LS), normalized least-squares (NLS), negative log-likelihood (NLL), and negative log-posterior (NLP) objectives, reported in Figure 3.18. All objective functions attain their minimum within the same region of the parameter space, demonstrating consistency between deterministic and probabilistic formulations. In particular, the NLL and NLP objectives preserve the location of the LS minimum while modifying the local geometry through the incorporation of uncertainty information.

The MAP estimate is given by

$$\Theta_{\text{MAP}} \approx (27.05, 34.02),$$

which is in close agreement with the reference configuration used to generate the target response. The corresponding predicted response is $\mu(\Theta_{\text{MAP}}) \approx (0.406, 0.381, 0.338)$, with predictive standard deviations of order 10^{-3} . The reconstruction error remains small, $\|\mathbf{y}_{\text{pred}} - \mathbf{y}^*\|_2 \approx 7.2 \times 10^{-4}$, confirming that the surrogate-based inversion accurately recovers both the target response and the associated parameter configuration.

The contour plots reveal a clear distinction between deterministic and uncertainty-aware objectives. The LS and NLS formulations exhibit relatively broad and nearly isotropic basins around the minimum, reflecting an unweighted data misfit. In contrast, the NLL and NLP objectives display sharper and anisotropic contours, induced by the weighting through the inverse total covariance matrix Σ_{tot} . This effect concentrates the objective landscape along directions of lower uncertainty and effectively penalizes regions associated with higher predictive variance or residual discrepancy.

No secondary minima are observed in the explored region, indicating that the inverse problem is locally well-posed. The local conditioning of the objective, characterized by a condition number close to unity (approximately 1.15), further supports stable identifiability in a neighborhood of the solution.

The effective covariance matrix Σ_{eff} is estimated from out-of-sample residuals and its structure is analyzed across calibration dataset sizes, as well as before and after adaptive refinement, as

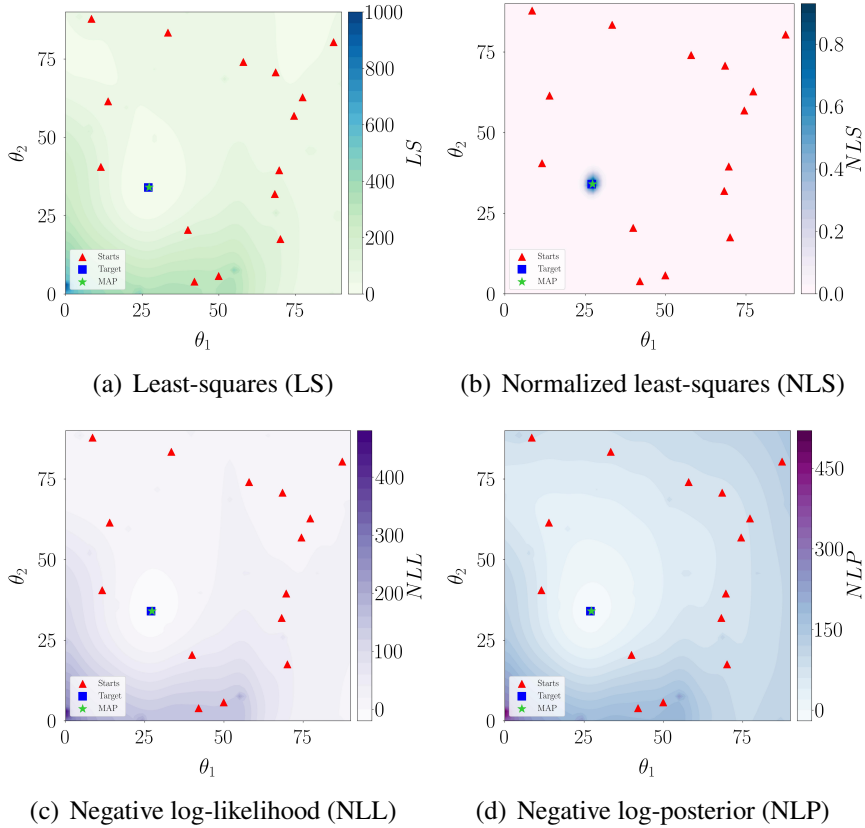


Figure 3.18: Contour plots of the objective functions over the parameter space (θ_1, θ_2) . Darker regions correspond to lower objective values. The blue square denotes the reference configuration, the green star indicates the MAP estimate, and the red triangles represent the initial multistart points. The comparison highlights the effect of uncertainty-aware weighting in shaping the objective landscape while preserving the location of the global minimum.

shown in Figure 3.19.

The estimated covariance matrices exhibit a predominantly diagonal structure in both the initial and refined configurations, indicating that residual discrepancies are largely uncorrelated across output components. This behavior is consistent with the independent multi-output Gaussian-process construction introduced in Section 3.2.4.1 and supports the separability assumption underlying the surrogate model.

The diagonal dominance persists across all calibration dataset sizes, while increasing the number of calibration samples improves the stability of the variance estimates. In particular, larger datasets reduce fluctuations in the diagonal entries without introducing significant off-diagonal contributions, suggesting that cross-correlations are intrinsically negligible rather than induced by limited sampling.

A comparison between the initial and refined covariance matrices shows a systematic

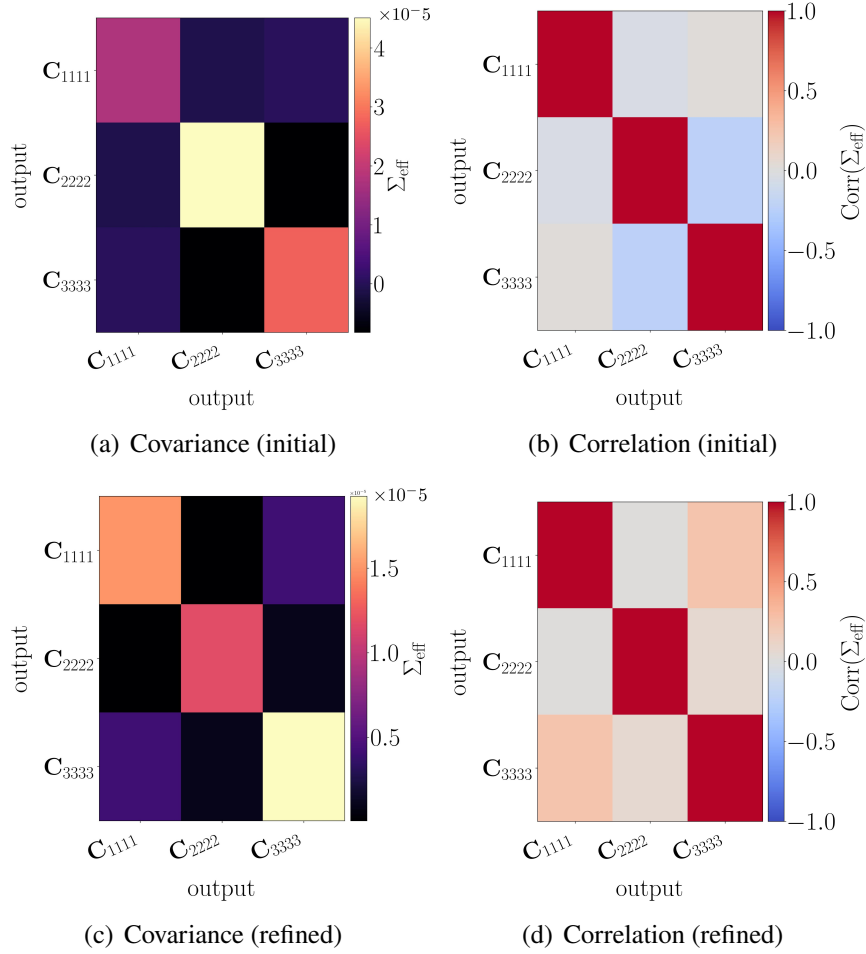


Figure 3.19: Structure of the effective covariance matrix estimated from residuals using a 20-sample calibration dataset. Top: covariance and correlation matrices prior to adaptive refinement. Bottom: corresponding matrices after posterior-focused refinement.

reduction in magnitude after adaptive refinement. This effect reflects a decrease in surrogate-model discrepancy in regions that contribute most to the posterior probability, where additional samples are concentrated by the refinement strategy. As a result, the residual uncertainty captured by Σ_{eff} becomes more localized and better aligned with the predictive variance of the surrogate model.

The corresponding correlation matrices confirm that refinement does not introduce artificial dependencies between output components. Off-diagonal terms remain close to zero in all cases, indicating that the structure of the residual uncertainty is preserved throughout the adaptive process.

The consistency of the MAP estimates across all calibration dataset sizes indicates that the inversion results are robust with respect to the estimation of Σ_{eff} . The inferred parameter

configurations remain confined to a small neighborhood of the reference solution, and no sensitivity to moderate variations in the covariance structure is observed. This behavior suggests that the likelihood formulation is well-conditioned and that the discrepancy term acts primarily as a regularization component without altering the location of the solution.

Overall, these results demonstrate that the proposed data-driven estimation of Σ_{eff} provides a stable and physically interpretable representation of residual uncertainty, and that its integration within the likelihood formulation leads to robust and consistent parameter inference.

The behavior of the adaptive refinement strategy is illustrated in Figure 3.20, where the evolution of the main components of the acquisition criterion is reported across refinement iterations.




Figure 3.20: Evolution of the components of the posterior-focused refinement criterion during adaptive learning (20-sample calibration case). The plots report, respectively, the predictive uncertainty at the selected points, the acquisition score, and the posterior weight driving the refinement strategy.

The selected predictive uncertainty exhibits a pronounced decrease during the first refinement iterations, indicating that the adaptive strategy rapidly reduces epistemic uncertainty in regions that are most relevant to the posterior distribution. After this initial phase, the uncertainty reduction becomes more gradual, suggesting that the surrogate model has already achieved a sufficiently accurate representation in the vicinity of the solution.

A consistent trend is observed for the refinement score, which combines uncertainty and posterior weighting. The score decreases sharply during the initial iterations and subsequently stabilizes, indicating that the acquisition process approaches a convergence regime where no regions simultaneously exhibit high uncertainty and high posterior relevance. This behavior provides evidence that the adaptive refinement effectively balances exploration and exploitation.

The evolution of the posterior weight further clarifies this transition. In the early iterations, relatively high values indicate that the refinement is driven by a broader exploration of the parameter space. As the iterations progress, the posterior weight becomes more concentrated, reflecting a progressive localization of the sampling process in a neighborhood of high posterior probability. This transition confirms that the refinement strategy successfully shifts from global exploration to posterior-focused exploitation.

The impact of bias correction is assessed through an ablation study comparing the standard Gaussian-process (GP) surrogate, the bias-corrected (BC) model, and direct inversion based on the high-fidelity model (HFM). Both GP and BC enable accurate parameter recovery with a limited number of high-fidelity evaluations, whereas direct HFM-based inversion does not converge within the same computational budget. The BC model consistently yields lower values of the negative log-likelihood and negative log-posterior, indicating improved agreement between the surrogate predictions and the underlying high-fidelity model, as well as enhanced internal consistency of the probabilistic formulation.

The inversion is performed through maximum a posteriori (MAP) estimation, resulting in a point estimate of the most probable parameter configuration. The observed objective landscape is unimodal and strongly concentrated, supporting the validity of MAP-based inference in the present setting. A full characterization of posterior uncertainty would require sampling-based techniques, which are not considered here.

The adaptive refinement strategy depends on the hyperparameters β_d , β_{AR} , and T , set to $\beta_d = 2$, $\beta_{AR} = 2$, and $T = 1$. Numerical evidence indicates that the inversion results remain stable under moderate variations of these parameters, suggesting limited sensitivity of the framework to their specific choice.

Overall, the results demonstrate that the proposed uncertainty-aware framework enables efficient and robust solution of the inverse design problem. The combination of surrogate modeling, discrepancy-aware likelihood formulation, and posterior-focused adaptive refinement leads to accurate parameter identification while maintaining computational efficiency.

3.2.5.2 Results in Four-Dimensional Parameter Space

The performance of the uncertainty-aware Bayesian inversion framework is assessed in a four-dimensional parameter space, where all parameters governing the spinodoid morphology are treated as unknown. The design vector is defined as $\Theta = (\theta_1, \theta_2, \theta_3, v_f)$, with admissible domain $\theta_1, \theta_2, \theta_3 \in [0^\circ, 90^\circ]$ and $v_f \in [0, 1]$, consistently with the parametrization introduced in Section 3.2.2.

Compared to the two-dimensional setting, the present configuration introduces increased complexity due to the higher dimensionality of the parameter space and the stronger coupling between geometric descriptors and effective properties. This setting therefore provides a more demanding test for surrogate modeling and Bayesian inversion in terms of sample efficiency, parameter identifiability, and stability of the likelihood formulation.

The inverse problem is defined in a synthetic setting by selecting a reference configuration $\Theta^* = (27^\circ, 34^\circ, 15^\circ, 0.45)$, from which the target response \mathbf{y}^* is generated through evaluation of the high-fidelity forward model \mathcal{F} (Equation (3.25)). As in the two-dimensional case, no observational noise is introduced and no additional model discrepancy is prescribed beyond the surrogate approximation, so that the contribution of surrogate modeling, predictive uncertainty, and data-driven discrepancy can be isolated within the likelihood formulation.

The surrogate model is constructed using the adaptive Gaussian-process procedure described in Section 3.2.4.1. An initial training set of 12 samples is progressively enriched over 80 adaptive

iterations using a selection criterion based on integrated predictive variance and a distance-based term in the normalized parameter space. Independent Gaussian-process models are employed for each output component using a Matérn kernel with smoothness parameter $\nu = 2.5$, with hyperparameters estimated by maximum likelihood.

Parameter inference is performed in a maximum a posteriori (MAP) setting following the framework introduced in Section 3.2.4.2. A multistart strategy with 15 initial points is adopted, combining basin-hopping with local L-BFGS-B optimization. The likelihood incorporates both surrogate predictive uncertainty and an effective covariance matrix Σ_{eff} , estimated from an independent calibration dataset of 200 samples not used for surrogate construction, leading to a total covariance $\Sigma_{\text{tot}} = \Sigma_{\text{eff}} + \Sigma_{\text{GP}}(\Theta)$ that accounts for both local predictive uncertainty and residual discrepancies. A uniform prior distribution is assumed over the admissible domain.

Posterior-focused adaptive refinement is performed over 5 iterations using the criterion defined in Equation (3.44), with the objective of improving surrogate accuracy in regions of high posterior probability. In addition, a bias-corrected surrogate model is introduced through a secondary Gaussian process trained on cross-validated residuals, enabling the effect of systematic discrepancy reduction to be assessed in the higher-dimensional setting.

3.2.5.2.1 Adaptive Construction of Gaussian-Process Surrogate Models

The performance of the Gaussian-process surrogate model in the four-dimensional parameter space is evaluated in terms of convergence behavior, predictive accuracy, and consistency with the underlying high-fidelity mapping.

The surrogate model approximates the forward relation defined in Equation (3.28), with output vector $\mathbf{y} = (C_{1111}^{\text{eff}}, C_{2222}^{\text{eff}}, C_{3333}^{\text{eff}})$. Consistently with the independent multi-output construction introduced in Section 3.2.4.1, the assessment focuses on global performance metrics that characterize the overall quality of the approximation across the admissible domain.

The convergence of the adaptive construction is assessed through validation error metrics. Figure 3.21 reports the evolution of the normalized root mean square error (NRMSE) and the coefficient of determination R^2 as functions of the number of training samples.

At the initial stage, the surrogate model exhibits relatively large errors, reflecting both the limited coverage of the parameter space provided by the initial design and the increased dimensionality of the problem. The activation of adaptive sampling, driven by the integrated predictive variance defined in Equation (3.35), leads to a rapid reduction of the approximation error during the first enrichment iterations. This initial regime is followed by a slower convergence phase, in which the NRMSE decreases more gradually and stabilizes at values of order 10^{-3} after approximately 40–50 samples, with only marginal improvements in subsequent iterations.

A consistent behavior is observed for the coefficient of determination. The R^2 metric rapidly increases during the early enrichment phase, exceeding 0.99 after a limited number of iterations and remaining close to unity throughout the remainder of the adaptive process. This trend indicates that the surrogate model captures the dominant functional dependence of the effective properties on the design parameters, despite the increased complexity of the mapping.

In contrast to the two-dimensional case, a component-wise or geometric visualization of the surrogate response is not reported. Due to the four-dimensional nature of the parameter space,

low-dimensional projections would provide only partial and potentially misleading representations of the underlying mapping, as they neglect higher-order interactions among parameters. For this reason, the assessment is based on global validation metrics, which provide a consistent and dimension-independent characterization of predictive performance.




Figure 3.21: Convergence of the Gaussian-process surrogate model during adaptive enrichment in the four-dimensional parameter space. The NRMSE exhibits a rapid initial decrease followed by gradual stabilization, while R^2 quickly approaches unity and remains stable throughout the enrichment process, confirming accurate reconstruction of the high-fidelity mapping.

Overall, the surrogate model achieves stable convergence and good predictive accuracy in the four-dimensional parameter space. Compared to the two-dimensional case, a reduction in convergence rate and a higher residual error are observed, which are consistent with the increased dimensionality and complexity of the underlying structure–property mapping. Despite these challenges, the adaptive Gaussian-process construction effectively captures the dominant features of the high-fidelity response and provides a sufficiently accurate and robust approximation for the subsequent inversion stage.

This level of accuracy is essential for the uncertainty-aware Bayesian inversion framework, as both the predictive mean and the predictive variance directly contribute to the likelihood formulation through Σ_{tot} , and therefore influence the inferred posterior distribution.

3.2.5.2.2 Uncertainty-Aware Bayesian Inversion Results

The performance of the uncertainty-aware Bayesian inversion framework in the four-dimensional parameter space is assessed by analyzing reconstruction accuracy, probabilistic consistency, and identifiability of the inverse problem.

The MAP estimate obtained using the bias-corrected (BC) surrogate model is given by

$$\Theta_{\text{MAP}} \approx (27.14, 34.09, 14.65, 0.406),$$

which lies in close proximity to the reference configuration

$$\Theta^* = (27^\circ, 34^\circ, 15^\circ, 0.45).$$

The reconstruction is highly accurate for the anisotropy angles, while a larger deviation is observed for the volume fraction. This behavior reflects the heterogeneous sensitivity of the effective response with respect to the individual parameters. In particular, the anisotropy angles directly influence the directional stiffness components and are therefore strongly informed by the observed response, whereas the volume fraction contributes in a more distributed manner, resulting in weaker identifiability.

The predicted response at the MAP configuration is

$$\boldsymbol{\mu}(\boldsymbol{\Theta}_{\text{MAP}}) \approx (0.321, 0.268, 0.272),$$

with predictive standard deviations of order 10^{-2} – 10^{-1} . The resulting data misfit,

$$\|\mathbf{y}_{\text{pred}} - \mathbf{y}^*\|_2 \approx 1.23 \times 10^{-1},$$

remains consistent with the uncertainty levels encoded in the likelihood formulation. This is further supported by the values of the negative log-posterior and negative log-likelihood,

$$\text{NLP}_{\text{BC}} \approx -3.65, \quad \text{NLL}_{\text{BC}} \approx -3.55,$$

indicating that the reconstructed solution lies within a region of high posterior probability. A comparison with the standard Gaussian-process (GP) surrogate model highlights the role of bias correction. The GP formulation yields a larger reconstruction error and higher objective values, indicating that residual systematic discrepancies in the surrogate approximation propagate into the inversion when not explicitly accounted for. The bias-corrected formulation mitigates this effect, leading to improved agreement with the target response and enhanced probabilistic consistency.

The structure of the effective covariance matrix $\boldsymbol{\Sigma}_{\text{eff}}$, reported in Figure 3.22, provides insight into the nature of the residual uncertainty.

The covariance matrices exhibit a predominantly diagonal structure, indicating that a large portion of the residual variance is associated with individual output components. However, non-negligible off-diagonal terms are present, revealing moderate correlations between the effective stiffness coefficients. These correlations reflect the coupled dependence of the homogenized response on the underlying microstructural parameters and support the use of a full covariance structure in the likelihood formulation. A comparison between the initial and refined covariance matrices shows a systematic reduction in magnitude after adaptive refinement, indicating that posterior-focused enrichment effectively reduces surrogate-model discrepancies in regions of high posterior probability. At the same time, the correlation structure remains qualitatively unchanged, suggesting that refinement improves local accuracy without altering the intrinsic dependency structure of the residual errors.

The behavior of the posterior-focused adaptive refinement strategy is illustrated in Figure 3.23.

The selected predictive uncertainty decreases across iterations, indicating a progressive reduction of epistemic uncertainty in posterior-relevant regions. The refinement score exhibits a rapid decrease followed by stabilization, suggesting convergence of the adaptive procedure.




Figure 3.22: Effective covariance matrix Σ_{eff} estimated from residuals in the four-dimensional setting. Top: covariance and correlation matrices before posterior-focused refinement. Bottom: corresponding matrices after refinement.

The posterior weight reflects a transition from exploratory sampling to concentration in regions of higher posterior probability, confirming that the refinement strategy effectively balances exploration and exploitation and progressively focuses the sampling process in regions that contribute most to the posterior distribution.

Despite the accurate reconstruction provided by the MAP estimate, the inverse problem exhibits pronounced anisotropy. The associated condition numbers are of order 10^{16} , indicating severe ill-conditioning and strong directional sensitivity of the objective landscape. The corresponding eigenvalue spectrum spans several orders of magnitude, revealing directions in the parameter space that are only weakly informed by the observed response. This behavior indicates that accurate pointwise reconstruction does not imply uniform identifiability across all parameters. Rather, the inverse problem combines good local recovery of the target configuration




Figure 3.23: Evolution of the components of the posterior-focused refinement criterion during adaptive learning.

with intrinsic non-identifiability along specific parameter combinations.

Direct inversion based on the high-fidelity forward model does not yield a valid solution within the same computational budget, confirming that brute-force approaches are computationally infeasible in the four-dimensional setting. In contrast, the surrogate-based framework enables efficient exploration of the parameter space and accurate recovery of the dominant features of the target configuration with a limited number of high-fidelity evaluations. Overall, the results demonstrate that the proposed uncertainty-aware Bayesian inversion framework provides an uncertainty-consistent reconstruction of the inverse design problem in higher-dimensional parameter spaces. The combination of surrogate modeling, discrepancy-aware likelihood formulation, and posterior-focused adaptive refinement enables stable and computationally efficient inference, while explicitly revealing the intrinsic identifiability limitations of the problem.

3.2.5.3 Conclusions

The numerical results presented in this section provide a systematic assessment of the uncertainty-aware Bayesian inversion framework across parameter spaces of increasing dimensionality. The analysis has focused on the interaction between surrogate modeling accuracy, probabilistic consistency of the likelihood formulation, and identifiability of the inverse problem.

In the two-dimensional setting, the adaptive Gaussian-process surrogate model achieves high predictive accuracy, as evidenced by the rapid convergence of validation metrics and the accurate reconstruction of the underlying structure–property mapping. The corresponding inversion results demonstrate that the proposed likelihood formulation preserves the geometry of the deterministic data misfit while incorporating uncertainty information in a consistent manner. The MAP estimates closely recover the reference configuration, and the objective landscape exhibits favorable conditioning and local identifiability. The effective covariance matrix remains predominantly diagonal, indicating that residual discrepancies are largely component-wise, and its estimation is stable with respect to the calibration dataset size. The posterior-focused adaptive

refinement strategy effectively concentrates sampling in regions of high posterior probability, leading to a reduction of epistemic uncertainty and improved local surrogate accuracy.

In the four-dimensional setting, the inverse design problem becomes significantly more challenging due to the increased dimensionality and stronger parameter interactions. The surrogate model maintains good predictive accuracy, although with a reduced convergence rate and higher residual error compared to the two-dimensional case. The inversion results remain consistent, with the MAP estimate accurately recovering the anisotropy parameters and providing a reasonable approximation of the volume fraction. The likelihood formulation remains probabilistically coherent, as the observed data misfit is compatible with the uncertainty levels encoded in the total covariance model. The effective covariance matrix reveals moderate cross-output correlations, reflecting the coupled nature of the effective mechanical response, while posterior-focused refinement continues to improve surrogate accuracy in regions of high posterior probability.

Despite the accurate reconstruction of the target configuration, the four-dimensional inverse problem exhibits pronounced anisotropy and severe ill-conditioning. The associated condition numbers indicate strong directional sensitivity of the objective landscape, with certain parameter combinations only weakly informed by the observed response. These results highlight an intrinsic limitation of the inverse design problem, namely that accurate pointwise reconstruction does not imply uniform identifiability across the parameter space. Rather, the inverse problem admits directions along which variations in the parameters produce limited changes in the observable quantities, leading to flat regions in the objective landscape and reduced parameter sensitivity.

A comparison between the two- and four-dimensional settings reveals a clear trade-off between reconstruction accuracy and identifiability. While the proposed framework remains effective in recovering the dominant features of the target configuration, increased dimensionality leads to reduced sensitivity and stronger non-identifiability along specific directions. This behavior is primarily governed by the structure–property mapping defined by the high-fidelity forward model and cannot be fully mitigated by surrogate modeling alone.

These observations suggest several directions for methodological extension. First, the identifiability limitations observed in the four-dimensional setting could be alleviated by enriching the observation space, for instance by incorporating additional components of the effective elasticity tensor or alternative mechanical descriptors, thereby increasing the sensitivity of the forward mapping to the design parameters. Second, the use of parameter-dependent discrepancy models or hierarchical Bayesian formulations could provide a more refined representation of residual model errors, improving the consistency of the likelihood in regions of sparse data. Third, sampling-based inference strategies, such as Markov chain Monte Carlo methods, could be employed to characterize the full posterior distribution and explicitly quantify uncertainty along poorly identifiable directions, rather than relying solely on point estimates. Finally, extensions toward multi-fidelity or multi-output correlated surrogate models may further improve sample efficiency and capture cross-output dependencies directly at the surrogate level.

Overall, the results demonstrate that the proposed framework provides a coherent and computationally efficient approach for surrogate-based Bayesian inversion in spinodoid architected materials. The combination of adaptive Gaussian-process surrogate modeling, discrepancy-aware

likelihood formulation, and posterior-focused refinement enables accurate and uncertainty-consistent parameter inference, while explicitly revealing the intrinsic limitations associated with high-dimensional inverse design problems and providing a basis for future methodological developments.

3.3 Conclusions

The present chapter examined surrogate-based methodologies for inverse problems within a Bayesian setting, with emphasis on the construction of computationally efficient surrogate models and on the explicit treatment of uncertainty in parameter inference. Two complementary contributions were presented, both based on the integration of adaptive Gaussian-process surrogate modeling and Bayesian inversion, but developed at different levels of generality and application complexity.

Section 3.1 introduced an uncertainty-aware framework for inverse problems in controlled low-dimensional benchmark settings. The analysis showed that adaptive Gaussian-process surrogate models, constructed through predictive-uncertainty-driven sampling, provide accurate approximations of the forward model with a limited number of high-fidelity evaluations. The resulting surrogate models enabled efficient Bayesian inversion, supporting parameter inference together with a local characterization of the posterior distribution through MAP estimation and Laplace approximation. The benchmark study clarified the interplay between surrogate accuracy, adaptive refinement, and posterior structure, and demonstrated the capability of the framework to capture both well-posed and multimodal inverse problem configurations.

Section 3.2 extended this surrogate-based Bayesian framework to the inverse design of spinoid architected materials, where the underlying structure–property mapping is nonlinear and anisotropic and the parameter space is of higher dimension. In this setting, the likelihood formulation incorporated both the predictive uncertainty of the Gaussian-process surrogate model and the residual discrepancy between surrogate predictions and high-fidelity evaluations through an effective covariance term. This formulation provided a more comprehensive probabilistic representation of epistemic uncertainty and supported stable parameter inference in reduced two- and four-dimensional design spaces.

Taken together, the two contributions show that adaptive surrogate modeling and Bayesian inversion can be combined within a coherent probabilistic framework for inverse problems involving computationally expensive forward models. The results indicate that the treatment of surrogate-induced uncertainty and model discrepancy influences the stability, identifiability, and interpretability of the inferred solutions.

The analyses were conducted in controlled settings, which enabled a systematic assessment of the effects of surrogate approximation, adaptive sampling, and likelihood modeling on the posterior distribution. This setting also highlights an important methodological aspect: the quality of parameter inference depends not only on the global accuracy of the surrogate model, but more specifically on its accuracy in the regions of the parameter space that govern the likelihood and posterior structure. For this reason, adaptive sampling strategies play a central

role not only in improving surrogate predictions, but also in supporting reliable and interpretable Bayesian inversion.

The scope of the present conclusions is therefore restricted to low-dimensional problems in which the interaction between surrogate modeling and Bayesian inference can be analyzed in a controlled manner. Extensions to more general scenarios, including experimental observations, heterogeneous noise models, stochastic forward operators, and higher-dimensional parameterizations, remain open research directions.

Overall, the chapter establishes a consistent methodological basis for uncertainty-aware inverse analysis based on Gaussian-process surrogate modeling and Bayesian inversion. The first contribution provides a detailed characterization of the framework on analytical benchmarks, while the second demonstrates its applicability to an inverse design problem of direct relevance in architected materials. This combination provides both methodological insight and application-oriented evidence supporting the use of surrogate-based Bayesian approaches in computational mechanics.

Chapter 4

Structural optimization strategies with optimality criteria and method of moving asymptotes on two- and three-dimensional benchmarks

4.1 Introduction

Structural topology optimization represents a central component of computational mechanics and provides a mathematically rigorous framework for the synthesis of optimal structural layouts under prescribed performance objectives and physical constraints [11, 31, 180, 218, 219]. In contrast to size and shape optimization, topology optimization allows changes in structural connectivity and enables the automatic generation of complex layouts through the distribution of material within a fixed design domain [220, 221]. This capability makes topology optimization particularly suitable for performance-driven design problems in aerospace, mechanical, and civil engineering, where stiffness, strength, and material efficiency must be addressed simultaneously [222].

The theoretical foundations of topology optimization originate from the seminal work of Bendsøe and Kikuchi on numerical homogenization [223, 224], which established a variational setting for material distribution problems governed by partial differential equations. Building upon this framework, topology optimization has developed into a broad research field encompassing a variety of formulations and numerical strategies, including density-based methods [32, 220], level-set approaches [225], topological derivatives [226], phase-field formulations [227], evolutionary algorithms [228], and hybrid techniques [229]. Among these methodologies, density-based approaches have gained widespread acceptance due to their conceptual simplicity, compatibility with standard finite element discretizations, and scalability to large-scale structural problems.

Within this class, the Solid Isotropic Material with Penalization (SIMP) method represents one of the most widely adopted formulations [31, 32]. The SIMP approach describes the design domain through a continuous material density field, where each point is modeled as an equivalent homogeneous material whose stiffness depends on the local density. Intermediate density values are penalized through a nonlinear interpolation law in order to promote nearly binary designs, consisting predominantly of solid and void regions, which is essential for manufacturability.

The structural response associated with a given density distribution is evaluated by solving the governing equilibrium equations using the Finite Element (FE) method [15, 16, 46, 53–57, 230], thereby coupling the optimization problem with a high-dimensional system of discretized partial differential equations.

Despite its practical success, the SIMP formulation leads to a nonconvex optimization problem characterized by a large number of design variables and strong coupling between the design and state equations. As a consequence, the numerical performance of SIMP-based topology optimization critically depends on the update strategy used to evolve the material density field during the iterative optimization process. A classical and computationally efficient approach in this context is provided by the Optimality Criterion (OC) method [31, 32, 231]. The OC method provides explicit update rules that balance objective reduction and constraint satisfaction, typically under a volume constraint [232]. Although conceptually simple, the interaction between penalization, filtering or regularization techniques, and OC-based updates gives rise to complex numerical behaviour, especially in nonconvex design landscapes and in the presence of active constraints. A rigorous analysis of both the continuum and discrete formulations is therefore essential to clarify the strengths and limitations of OC-based SIMP optimization.

Beyond OC strategies, structural topology optimization increasingly relies on general gradient-based methods capable of handling nonlinear objectives and multiple inequality constraints [233, 234]. Such methods typically construct local approximations of the objective and constraint functions based on first-order sensitivity information. However, in highly nonlinear or nonconvex problems, conventional gradient-based approaches may exhibit slow convergence, oscillatory behaviour, or numerical instability [235, 236]. These limitations motivate the development and use of optimization algorithms that incorporate mechanisms to control curvature and step size while preserving computational efficiency.

The Method of Moving Asymptotes (MMA), originally proposed by Svanberg [4], was developed to address these challenges. MMA is a first-order optimization method tailored for large-scale constrained problems, in which the original nonlinear and potentially nonconvex problem is replaced by a sequence of strictly convex and separable subproblems. A defining feature of the method is the introduction of adaptively updated lower and upper bounds, known as moving asymptotes, which regulate the curvature of the local approximations and implicitly control the admissible step size. Through this mechanism, MMA enhances numerical robustness and improves convergence behaviour in large-scale and strongly nonlinear optimization problems [33, 237–239].

Owing to these properties, MMA has become a reference method in structural and topology optimization and has been successfully applied to a wide range of problems, including minimum compliance optimization [34], large-scale topology optimization [240], stress-constrained formulations [241], and multidisciplinary design applications [242–244]. Despite its extensive use, the theoretical understanding of MMA remains only partially consolidated. Most contributions focus on algorithmic descriptions and implementation aspects, while a systematic analysis of the mathematical structure of the method — including convexity, separability, and first-order consistency — is often limited. This aspect hinders a rigorous interpretation of the method and obscures the connection between its analytical formulation and its observed numerical

performance.

The present chapter examines structural topology optimization through two complementary and widely used solution strategies. The chapter discusses two distinct research works developed during the doctoral programme, both currently in the pre-submission phase. The first contribution addresses topology optimization via the SIMP method combined with the OC update strategy, with particular emphasis on the continuum and discrete formulations, regularization techniques, numerical implementation, and benchmark analyses. The second contribution provides a comprehensive theoretical and algorithmic investigation of the MMA, revisiting its mathematical foundations, clarifying the role of local and separable convex approximations, and analysing the algorithmic mechanisms that govern convergence and numerical stability.

Although the methods considered in this chapter are well established in the literature, their analysis is conducted with the objective of assessing their suitability within a broader uncertainty-aware design framework. In particular, the insights gained from this study provide a basis for future developments aimed at integrating structural optimization with surrogate modeling and Bayesian inversion techniques.

Accordingly, Section 4.2 presents the SIMP-based topology optimization framework with OC updates, covering the governing equations, optimization formulation, numerical implementation, and representative benchmark problems. Section 4.3 is devoted to the Method of Moving Asymptotes and develops a rigorous theoretical and algorithmic analysis of the method, from the formulation of constrained structural optimization problems to the construction of convex approximations, the adaptive update of the moving asymptotes, and numerical validation. The chapter concludes in Section 4.4 with a synthesis of the main findings and a discussion of perspectives for advanced structural topology optimization.

4.2 Topology Optimization via the Solid Isotropic Material Penalization Method and an Optimality Criterion Approach

The present section presents the formulation and numerical treatment of structural topology optimization based on the SIMP method combined with an OC update strategy. The objective is to provide a coherent and rigorous description of the SIMP framework, starting from the continuum formulation of the governing equations Section 4.2.1, and proceeding through the corresponding discrete formulation Section 4.2.2, numerical implementation Section 4.2.3, representative benchmark examples Section 4.2.4, and discussing the results Section 4.2.5.

4.2.1 Continuum Formulation

In the present section, we discuss the continuum formulation of the topology optimization problem using the SIMP method, focusing on minimizing structural compliance [31, 32, 223, 224].

The primary design variable in SIMP method is the material density, denoted as ϕ , which is defined over the entire design domain $\Omega \subset \mathbb{R}^2$ (or \mathbb{R}^3) and varies within the interval $[0, 1]$

where $\phi = 0$ indicates void regions, while $\phi = 1$ represents solid regions. Intermediate values of ϕ , where $0 < \phi < 1$, represent transitional regions between solid and void states, modeling a continuous material distribution across the design domain.

The goal in topology optimization is to determine the optimal material distribution, defined by ϕ , across Ω , such that the structure meets specified mechanical performance requirements while minimizing material usage; in other words, the optimization aims to maximize structural stiffness while minimizing material use.

A key component in the SIMP method is the relationship between the design variable ϕ and the structural stiffness, which is described at each point in Ω by a penalized stiffness tensor dependent on ϕ . The penalized stiffness tensor, $\mathbb{C}(\phi)$, governs the relationship among applied loads, displacements, and material properties, thereby controlling the mechanical behavior of the structure. Specifically, $\mathbb{C}(\phi)$ is formulated as:

$$\mathbb{C}(\phi) = \phi^p \mathbb{C}_0, \quad (4.1)$$

where \mathbb{C}_0 denotes the stiffness tensor of the solid material, and p is the penalization factor, which penalizes intermediate density values to drive the solution toward a binary distribution (solid or void).

The penalized stiffness tensor $\mathbb{C}(\phi)$ plays a central role in determining the structural response, as it directly affects the relationship between external forces and resulting displacements. By influencing the stiffness of the structure, $\mathbb{C}(\phi)$ effectively guides the optimization process in adjusting the material distribution ϕ to achieve minimum compliance under specified load conditions while satisfying equilibrium constraints.

The remainder of the present section is organized as follows. Section 4.2.1.1 introduces the linear elastic problem, describing the mechanical behavior of the structure under load. Section 4.2.1.2 presents the formulation of the minimum compliance problem, defining the optimization objective. In Section 4.2.1.3, we derive the Lagrangian formulation, which serves as the basis for the optimization algorithm. Section 4.2.1.4 discusses the optimality criterion employed to update the material distribution. Finally, Section 4.2.1.5 outlines regularization techniques that ensure the stability of the solution.

4.2.1.1 Linear Elastic Problem

A linear elastic solid body occupying the design domain Ω is considered. The mechanical response is governed by the equilibrium, constitutive, and compatibility equations, supplemented by Dirichlet and Neumann boundary conditions. The strong form of the linear elastic problem

reads:

$$\begin{cases} \operatorname{div}(\boldsymbol{\sigma}) + \phi \mathbf{b} = \mathbf{0} & \text{in } \Omega, \\ \boldsymbol{\sigma} = \mathbb{C}(\phi) : \boldsymbol{\varepsilon} & \text{in } \Omega, \\ \boldsymbol{\varepsilon} = \nabla^s \mathbf{u} = \frac{\nabla \mathbf{u} + \nabla \mathbf{u}^T}{2} & \text{in } \Omega, \\ \mathbf{u} = \bar{\mathbf{u}} & \text{on } \Gamma_D, \\ \boldsymbol{\sigma} \mathbf{n} = \mathbf{t} & \text{on } \Gamma_N, \end{cases} \quad (4.2)$$

where $\boldsymbol{\sigma}$ denotes the Cauchy stress tensor, $\boldsymbol{\varepsilon}$ the infinitesimal strain tensor, and \mathbf{u} the displacement field. The vector \mathbf{b} represents the body force per unit volume, scaled by the material density field ϕ . The constitutive relation is governed by the density-dependent stiffness tensor $\mathbb{C}(\phi)$, defined according to the SIMP interpolation law introduced in Equation (4.1). Prescribed displacements $\bar{\mathbf{u}}$ are enforced on the Dirichlet boundary Γ_D , while surface tractions \mathbf{t} are applied on the Neumann boundary Γ_N .

For numerical discretization, the equilibrium equation is expressed in weak form. The weak formulation consists in finding a displacement field \mathbf{u} such that:

$$\int_{\Omega} \nabla^s \mathbf{u} : \mathbb{C}(\phi) : \nabla^s \delta \mathbf{u} \, d\Omega - \int_{\Omega} \phi \mathbf{b} \cdot \delta \mathbf{u} \, d\Omega - \int_{\Gamma_N} \mathbf{t} \cdot \delta \mathbf{u} \, d\Gamma = 0 \quad \forall \delta \mathbf{u}, \quad (4.3)$$

where $\delta \mathbf{u}$ denotes an admissible virtual displacement field vanishing on Γ_D .

4.2.1.2 Minimum Compliance Problem

In the present section, we address the topology optimization process for minimizing structural compliance within the SIMP framework. The compliance is represented by the functional $C(\phi, \mathbf{u})$, defined as:

$$C(\phi, \mathbf{u}) = \int_{\Omega} \phi \mathbf{b} \cdot \mathbf{u} \, d\Omega + \int_{\Gamma_N} \mathbf{t} \cdot \mathbf{u} \, d\Gamma, \quad (4.4)$$

The design variable ϕ influences the penalized stiffness tensor $\mathbb{C}(\phi)$ as defined in Equation (4.1), which in turn affects the equilibrium displacements $\mathbf{u}(\phi)$. Therefore, compliance $C(\phi, \mathbf{u})$ implicitly depends on both ϕ and \mathbf{u} .

To control material usage, a volume constraint is imposed, limiting the total material volume within the optimized design to a specified fraction f of the initial design domain volume, V . The volume constraint is formulated as:

$$\mathcal{V}(\phi) = \int_{\Omega} \phi \, d\Omega - fV \leq 0, \quad (4.5)$$

where $f \in (0, 1]$ denotes the allowable fraction of material.

In summary, the SIMP-based topology optimization problem is formulated as a constrained optimization problem, where the objective is to minimize compliance (Equation (4.4)) subject

to the equilibrium condition (Equation (4.3)) and the volume constraint (Equation (4.5)). The equilibrium condition, expressed in its weak form (Equation (4.3)), is here rewritten compactly as $S(\phi, \mathbf{u}, \delta \mathbf{p}) = 0$. This notation captures the requirement that the structure is in mechanical equilibrium for a given material distribution ϕ and corresponding displacement field \mathbf{u} . Together with the volume constraint (Equation (4.5)), which restricts the amount of material used, the optimization problem is expressed as:

$$\left\{ \begin{array}{ll} \text{minimize:} & C(\phi, \mathbf{u}), \\ \text{subject to:} & S(\phi, \mathbf{u}, \delta \mathbf{p}) = 0, \\ \text{with:} & \mathcal{V}(\phi) \leq 0, \\ \text{and:} & 0 \leq \phi \leq 1, \end{array} \right. \quad (4.6)$$

This formulation provides a rigorous framework for determining the optimal material distribution ϕ that satisfies both structural performance requirements and material constraints within the SIMP method.

4.2.1.3 Lagrangian Formulation

To address the minimum compliance problem, we employ a Lagrangian formulation that simultaneously incorporates the volume constraint and the equilibrium condition. In this approach, a Lagrange multiplier λ is introduced to enforce the volume constraint, as described in Equation (4.5). The Lagrangian formulation defines the optimization problem as the stationarity of a Lagrangian functional, which combines the compliance, volume constraint, and equilibrium condition. The Lagrangian functional is defined as:

$$\mathcal{L}(\phi, \mathbf{u}, \lambda, \mathbf{p}) = C(\phi, \mathbf{u}) + \lambda \mathcal{V}(\phi) + S(\phi, \mathbf{u}, \mathbf{p}), \quad (4.7)$$

where the compliance functional $C(\phi, \mathbf{u})$ is given by Equation (4.4) and the volume constraint $\mathcal{V}(\phi)$ is defined as in Equation (4.5). The term $S(\phi, \mathbf{u}, \mathbf{p})$ enforces the equilibrium condition and is expressed as:

$$S(\phi, \mathbf{u}, \mathbf{p}) = \int_{\Omega} \nabla^s \mathbf{u} : \mathbb{C}(\phi) : \nabla^s \mathbf{p} \, d\Omega - \int_{\Omega} \phi \mathbf{b} \cdot \mathbf{p} \, d\Omega - \int_{\Gamma_N} \mathbf{t} \cdot \mathbf{p} \, d\Gamma, \quad (4.8)$$

where $\mathbb{C}(\phi)$ is the penalized stiffness tensor, which depends on the design variable ϕ as prescribed in Equation (4.1).

The use of \mathbf{p} is particularly significant because it enables efficient sensitivity analysis of the compliance functional $C(\phi, \mathbf{u})$ with respect to the design variable ϕ . The compliance functional depends on ϕ and \mathbf{u} , with \mathbf{u} depending implicitly on ϕ through the equilibrium equation. A direct computation of the derivative $\frac{\partial C}{\partial \phi}$ would require evaluating $\frac{\partial \mathbf{u}}{\partial \phi}$, a process that is computationally expensive. To avoid this challenge, \mathbf{p} is introduced as an auxiliary variable associated with the equilibrium equation, allowing the sensitivity analysis to be reformulated without explicitly computing $\frac{\partial \mathbf{u}}{\partial \phi}$. This approach is called the adjoint method and \mathbf{p} is referred to as the adjoint variable.

Expanding the Lagrangian functional yields the following expression:

$$\begin{aligned} \mathcal{L}(\phi, \mathbf{u}, \lambda, \mathbf{p}) = & \int_{\Omega} \phi \mathbf{b} \cdot \mathbf{u} \, d\Omega + \int_{\Gamma_N} \mathbf{t} \cdot \mathbf{u} \, d\Gamma + \lambda \left(\int_{\Omega} \phi \, d\Omega - fV \right) \\ & + \int_{\Omega} \nabla^s \mathbf{u} : \mathbb{C}(\phi) : \nabla^s \mathbf{p} \, d\Omega - \int_{\Omega} \phi \mathbf{b} \cdot \mathbf{p} \, d\Omega - \int_{\Gamma_N} \mathbf{t} \cdot \mathbf{p} \, d\Gamma. \end{aligned} \quad (4.9)$$

The stationarity of the Lagrangian functional is achieved when its variation vanishes for any admissible variations in ϕ , \mathbf{u} , λ , and \mathbf{p} :

$$d\mathcal{L}(\phi, \mathbf{u}, \lambda, \mathbf{p}) [\delta\phi, \delta\mathbf{u}, \delta\lambda, \delta\mathbf{p}] = 0 \quad \forall \delta\phi, \delta\mathbf{u}, \delta\lambda, \delta\mathbf{p}. \quad (4.10)$$

This stationarity condition decomposes into four equations:

$$d\mathcal{L}(\phi, \mathbf{u}, \lambda, \mathbf{p}) [\delta\phi, 0, 0, 0] = 0, \quad (4.11a)$$

$$d\mathcal{L}(\phi, \mathbf{u}, \lambda, \mathbf{p}) [0, \delta\mathbf{u}, 0, 0] = 0, \quad (4.11b)$$

$$d\mathcal{L}(\phi, \mathbf{u}, \lambda, \mathbf{p}) [0, 0, \delta\lambda, 0] = 0, \quad (4.11c)$$

$$d\mathcal{L}(\phi, \mathbf{u}, \lambda, \mathbf{p}) [0, 0, 0, \delta\mathbf{p}] = 0. \quad (4.11d)$$

The explicit equations corresponding to each condition are as follows:

1. Stationarity with respect to ϕ (Equation (4.11a)):

$$\int_{\Omega} \delta\phi \mathbf{b} \cdot \mathbf{u} \, d\Omega + \lambda \int_{\Omega} \delta\phi \, d\Omega + \int_{\Omega} \nabla^s \mathbf{u} : \frac{\partial \mathbb{C}(\phi)}{\partial \phi} \delta\phi : \nabla^s \mathbf{p} \, d\Omega - \int_{\Omega} \delta\phi \mathbf{b} \cdot \mathbf{p} \, d\Omega = 0. \quad (4.12)$$

2. Stationarity with respect to \mathbf{u} (Equation (4.11b)):

$$\int_{\Omega} \nabla^s \delta\mathbf{u} : \mathbb{C}(\phi) : \nabla^s \mathbf{p} \, d\Omega - \int_{\Omega} \phi \mathbf{b} \cdot \delta\mathbf{u} \, d\Omega - \int_{\Gamma_N} \mathbf{t} \cdot \delta\mathbf{u} \, d\Gamma = 0. \quad (4.13)$$

This equation defines the adjoint problem, crucial for resolving the dependence of \mathbf{u} on ϕ in the sensitivity analysis.

3. Stationarity with respect to λ (Equation (4.11c)):

$$\int_{\Omega} \phi \, d\Omega - fV = 0. \quad (4.14)$$

This enforces the volume constraint.

4. Stationarity with respect to \mathbf{p} (Equation (4.11d)):

$$\int_{\Omega} \nabla^s \mathbf{u} : \mathbb{C}(\phi) : \nabla^s \delta\mathbf{p} \, d\Omega - \int_{\Omega} \phi \mathbf{b} \cdot \delta\mathbf{p} \, d\Omega - \int_{\Gamma_N} \mathbf{t} \cdot \delta\mathbf{p} \, d\Gamma = 0. \quad (4.15)$$

This equation corresponds to the equilibrium condition, ensuring mechanical balance.

The adjoint and equilibrium problems (Equation (4.13) and Equation (4.15)) are structurally identical due to the symmetry of $\mathbb{C}(\phi)$, a consequence of the linear elastic problem definition. This symmetry implies $\mathbf{p} = \mathbf{u}$, and the problem is called self-adjoint.

Substituting $\mathbf{p} = \mathbf{u}$ into the Lagrangian functional, it simplifies to depend only on ϕ and λ :

$$\mathcal{L}_{\text{eq}}(\phi, \lambda) = C_{\text{eq}}(\phi) + \lambda \mathcal{V}(\phi), \quad (4.16)$$

where $C_{\text{eq}}(\phi)$ represents the compliance functional under equilibrium conditions. This reduced form of the Lagrangian functional is referred to as the equilibrated Lagrangian functional, as it explicitly incorporates the equilibrium condition.

The term “equilibrated” emphasizes that the equilibrium condition is satisfied and fully integrated into the functional. As a result, the optimization process is simplified and focuses on determining the material distribution ϕ that minimizes compliance and respects volume constraints.

The reduced optimization problem is then solved by finding the stationary conditions of \mathcal{L}_{eq} respect to ϕ and λ :

$$\frac{\partial \mathcal{L}_{\text{eq}}}{\partial \phi} = 0, \quad (4.17a)$$

$$\frac{\partial \mathcal{L}_{\text{eq}}}{\partial \lambda} = 0. \quad (4.17b)$$

These conditions provide the optimality criteria for the design variable ϕ , ensuring compliance minimization under the equilibrium and volume constraints.

4.2.1.4 Optimality Criterion

The Optimality Criterion (OC) [231] describes the conditions under which the material distribution within the design domain achieves compliance minimization while satisfying the volume constraint. This criterion originates from the stationarity of the equilibrated Lagrangian functional \mathcal{L}_{eq} , as defined in Section 4.2.1.3.

The stationarity of \mathcal{L}_{eq} with respect to the design variable ϕ provides the necessary condition:

$$\frac{\partial \mathcal{L}_{\text{eq}}(\phi, \lambda)}{\partial \phi} = \frac{\partial C_{\text{eq}}(\phi)}{\partial \phi} + \lambda \frac{\partial \mathcal{V}(\phi)}{\partial \phi} = 0. \quad (4.18)$$

Rewriting this condition, the relationship between the compliance sensitivity and the volume constraint sensitivity takes the form:

$$-\frac{\partial C_{\text{eq}}(\phi)}{\partial \phi} \left(\lambda \frac{\partial \mathcal{V}(\phi)}{\partial \phi} \right)^{-1} = 1. \quad (4.19)$$

This equality constitutes the optimality condition, indicating that the material distribution ϕ balances the sensitivity of compliance with the sensitivity of the volume constraint at every point in the domain.

However, during the optimization process, this condition is not necessarily satisfied exactly at all points in the domain. To quantify the deviation from the optimality condition, the control variable B_{eq} is introduced:

$$B_{eq} = -\frac{\partial C_{eq}(\phi)}{\partial \phi} \left(\lambda \frac{\partial \mathcal{V}(\phi)}{\partial \phi} \right)^{-1}. \quad (4.20)$$

The variable B_{eq} serves as a measure of how closely the current material distribution aligns with the theoretical optimality condition. When $B_{eq} = 1$ throughout the domain, the compliance is minimized, and the volume constraint is satisfied, signifying that the material distribution is optimal.

In practice, the optimization process iteratively updates ϕ to progressively satisfy the optimality condition. Deviations of B_{eq} from unity during intermediate iterations highlight areas where the material distribution requires adjustment. By incorporating sensitivity analysis, the iterative procedure ensures convergence toward a material distribution where $B_{eq} \rightarrow 1$, achieving compliance minimization while respecting the volume constraint. Details of the iterative algorithm and its numerical implementation are discussed in Section 4.2.2.

4.2.1.5 Regularization Techniques

In topology optimization, regularization techniques [32, 220, 245, 245–248] are essential for obtaining physically meaningful and stable solutions. Without regularization, solutions can exhibit undesirable characteristics such as non-physical patterns, highly localized material concentrations, or abrupt transitions between solid and void regions [249, 250]. These phenomena compromise design stability and feasibility, highlighting the importance strategies that promote smoothness and stability in the material distribution.

Several regularization strategies are commonly used in topology optimization [32, 220, 245, 245–248], each addressing specific challenges: (i) filtering techniques [32, 220, 245], which smooth the material distribution by correlating the densities at neighbouring points in the design domain; (ii) perimeter control constraints [247], which limit the interface length of material regions to prevent overly complex designs; (iii) constraints on the slope of the density gradient [248], limiting rapid changes in material density across the domain; and (iv) direct constraints on the density gradient [249], promoting gradual transitions between material phases.

Among these, filtering methods [32, 220, 245] are widely employed due to their effectiveness in smoothing material transitions. The filtering process modifies the design variable ϕ by introducing a filter function \mathcal{K} , which adjusts the material density at each point in the domain based on the densities of neighbouring points. This regularization helps produce a stable material distribution and avoids non-physical transitions.

Filtering can be applied to the material density itself. The general form of a density-based filtering technique can be expressed as:

$$\phi(\mathbf{x}) = \mathcal{K}(\phi(\mathbf{y})), \quad (4.21)$$

where \mathbf{x} denotes a point in the design domain, and \mathbf{y} represents neighbouring points. The filter function \mathcal{K} averages the densities around each point \mathbf{x} , resulting in a smooth material distribution $\phi(\mathbf{x})$.

Additionally, filtering can be applied to the sensitivity of the objective function, specifically the sensitivity of the compliance at equilibrium C_{eq} , with respect to the material density. This approach, known as sensitivity-based filtering, further stabilizes the optimization process by smoothing the sensitivity values. Sensitivity-based filtering is formulated as:

$$\frac{\partial C_{\text{eq}}(\phi(\mathbf{x}))}{\partial \phi(\mathbf{x})} = \mathcal{K} \left(\frac{\partial C_{\text{eq}}(\phi(\mathbf{y}))}{\partial \phi(\mathbf{y})} \right). \quad (4.22)$$

This formulation helps ensure stable optimization steps and prevents sharp local variations in material distribution.

A commonly used type of filter relies on Radial Basis Functions (RBF) [246], which smooth the material distribution based on the distance between points in the design domain. RBF filters are compactly supported, meaning they influence points only within a defined radius, ensuring localized smoothing. The RBF-based filter is given by:

$$\mathcal{K} = \kappa(r(\mathbf{x}, \mathbf{y})) = [r_0(\mathbf{x}, \mathbf{y})]^4 (4r(\mathbf{x}, \mathbf{y}) + 1), \quad (4.23)$$

where $r(\mathbf{x}, \mathbf{y}) = \frac{\|\mathbf{y}-\mathbf{x}\|}{r_{\text{min}}}$ is the normalized distance between points \mathbf{x} and \mathbf{y} , and r_{min} specifies the filter's radius of influence. The function $r_0(\mathbf{x}, \mathbf{y}) = \max(0, 1 - r(\mathbf{x}, \mathbf{y}))$ ensures compact support, limiting the filter's effect to points within a certain radius.

Using the RBF filter, the regularized material density can be defined as:

$$\phi(\mathbf{x}) = \int_{\Omega} \kappa(r(\mathbf{x}, \mathbf{y})) \phi(\mathbf{y}) d\Omega, \quad (4.24)$$

and the filtered sensitivity is given by:

$$\frac{\partial C_{\text{eq}}(\phi(\mathbf{x}))}{\partial \phi(\mathbf{x})} = \int_{\Omega} \kappa(r(\mathbf{x}, \mathbf{y})) \frac{\partial C_{\text{eq}}(\phi(\mathbf{y}))}{\partial \phi(\mathbf{y})} d\Omega. \quad (4.25)$$

While these regularization techniques establish a theoretical foundation for stable solutions in the continuum formulation, their practical importance is particularly evident in the discrete formulation. A detailed discussion on these techniques is presented in Section 4.2.2.5.

4.2.2 Discrete Formulation

In the present section, the continuum formulation is discretized in order to obtain a finite-dimensional optimization problem suitable for numerical solution using the Finite Element method.

As discussed in Section 4.2.1, the SIMP method assumes that regions with $\phi = 0$ correspond to void, while regions with $\phi = 1$ correspond to solid material. In the discrete setting, assigning

zero stiffness to void regions would lead to singular stiffness matrices. To avoid this issue, void regions are modeled as a very soft material by introducing a lower bound $\phi_{\min} > 0$. Accordingly, the penalized stiffness tensor is interpolated as:

$$\mathbb{C}(\phi) = [\phi_{\min} + (1 - \phi_{\min})\phi^p] \mathbb{C}_0, \quad (4.26)$$

where \mathbb{C}_0 denotes the stiffness tensor of the solid material, ϕ_{\min} is a small positive parameter introduced to ensure numerical stability, and p is the penalization exponent [31, 32].

The choice of the penalization exponent p plays a central role in the SIMP method, as it controls the tendency of the material distribution to approach a binary (0–1) design. For $p = 1$, the interpolation is linear and does not penalize intermediate densities, resulting in solutions characterized by significant regions with $0 < \phi < 1$, often referred to as “gray” regions. These intermediate densities are generally undesirable from a manufacturing perspective [31, 32].

To promote nearly binary designs, values of $p > 1$ are typically employed. In particular, values in the range $p \approx 3$ are commonly adopted in practical applications, as they provide an effective compromise between penalization of intermediate densities and numerical stability of the optimization process [31, 32, 34]. For such values, intermediate densities are strongly penalized, driving the solution towards a clear separation between solid and void regions.

Although larger values of p can further enhance the penalization effect, excessively high values may lead to ill-conditioned optimization problems, potentially slowing convergence or inducing numerical instabilities [32]. For this reason, moderate values of the penalization exponent are generally preferred in practice.

4.2.2.1 Linear Elastic Problem

In the present section, we introduce the discrete formulation of the linear elastic equilibrium problem within the SIMP-based topology optimization framework. Following the continuum formulation in Section 4.2.1, this discrete approach establishes equilibrium across a finite element discretization of the design domain, Ω , while retaining the constitutive and compatibility equations in their strong forms.

The design domain Ω is discretized into a finite number of elements, indexed by e , using a finite element mesh [230]. Within each element, the material density ϕ is considered constant, and the displacement field \mathbf{u} is interpolated at the element nodes using shape functions. Let \mathbf{N}_u represent the shape functions associated with the displacement field; then, the displacement \mathbf{u} within each element e can be approximated as:

$$\mathbf{u}_e \approx \mathbf{N}_u \mathbf{d}_e, \quad (4.27)$$

where \mathbf{d}_e is the vector of nodal displacements for element e .

The discrete formulation of the equilibrium problem requires defining the relationship between the Cauchy stress tensor $\boldsymbol{\sigma}_e$, the material stiffness tensor $\mathbb{C}(\phi_e)$, and the strain tensor $\boldsymbol{\varepsilon}_e$ within each element. The stress-strain relationship for element e is given by the constitutive law:

$$\boldsymbol{\sigma}_e = \mathbb{C}(\phi_e) : \boldsymbol{\varepsilon}_e, \quad (4.28)$$

where $\mathbb{C}(\phi_e)$ is the penalized stiffness tensor, which depends on the element material density ϕ_e as defined in the SIMP interpolation law (Equation (4.26)). The strain tensor $\boldsymbol{\varepsilon}_e$ is derived from the nodal displacements \mathbf{d}_e through the matrix of shape function derivatives \mathbf{B}_e as:

$$\boldsymbol{\varepsilon}_e = \mathbf{B}_e \mathbf{d}_e. \quad (4.29)$$

At the element level, the principle of virtual work is applied, yielding the element stiffness matrix $\mathbf{K}_e(\phi_e)$ as:

$$\mathbf{K}_e(\phi_e) = \int_{\Omega_e} \mathbf{B}_e^T \mathbb{C}(\phi_e) \mathbf{B}_e d\Omega_e, \quad (4.30)$$

where the integration is carried out over the element domain Ω_e . Here, the material stiffness tensor $\mathbb{C}(\phi_e)$ adjusts the stiffness of each element based on ϕ_e , transitioning from solid to void regions and ensuring the physical material distribution.

The global equilibrium equation is formed by assembling the contributions from all elements and is expressed as:

$$\mathbf{K} \mathbf{d} = \mathbf{f}, \quad (4.31)$$

where $\mathbf{K} = \mathbf{K}(\phi)$ is the global stiffness matrix, constructed by assembling the element stiffness matrices $\mathbf{K}_e(\phi_e)$ for all elements e , $\mathbf{d} = \mathbf{d}(\phi)$ is the global displacement vector, containing the nodal displacements for the entire finite element mesh and \mathbf{f} is the global force vector, incorporating both external and body forces acting on the structure.

For a given material distribution ϕ , the global equilibrium equation (Equation (4.31)) is solved to determine the nodal displacement vector \mathbf{d} . Solving the linear system in (Equation (4.31)) provides the necessary displacements to evaluate the structural response corresponding to the current material configuration.

In summary, the discrete formulation of the linear elastic equilibrium problem provides the basis for computing the displacement field \mathbf{d} for a given material distribution ϕ .

4.2.2.2 Minimum Compliance Problem

In the discrete setting of topology optimization, the goal is to minimize the structural compliance defined as:

$$c(\phi, \mathbf{d}) = \mathbf{f}^T \mathbf{d}, \quad (4.32)$$

with $\mathbf{d} = \mathbf{d}(\phi)$.

The volume constraint ensures that the total material used does not exceed a specified fraction of the initial design domain volume. This constraint can be expressed as:

$$v(\phi) = \sum_{e=1}^{n_e} \phi_e V_e - f V_0 \leq 0, \quad (4.33)$$

where V_e denotes the volume of element e and $f \in (0, 1]$ represents the allowable fraction of the initial design volume V_0 .

Therefore, the discrete minimum compliance problem is formulated as:

$$\begin{cases} \text{minimize:} & c(\phi, \mathbf{d}), \\ \text{subject to:} & \mathbf{K}\mathbf{d} = \mathbf{f}, \\ \text{with:} & v(\phi) \leq 0, \\ \text{and:} & 0 \leq \phi \leq 1. \end{cases} \quad (4.34)$$

4.2.2.3 Lagrangian Formulation

The equilibrated Lagrangian functional for the discrete problem is defined as

$$\mathcal{L}_{\text{eq}}(\phi, \lambda) = c_{\text{eq}}(\phi) + \lambda v(\phi), \quad (4.35)$$

where $c_{\text{eq}}(\phi)$ denotes the compliance at equilibrium for a given material distribution ϕ , and λ is the Lagrange multiplier associated with the volume constraint $v(\phi)$.

To derive the optimality condition, the sensitivity of c_{eq} with respect to ϕ is required. The discrete equilibrium equation reads $\mathbf{K}(\phi)\mathbf{d}(\phi) = \mathbf{f}$ (cf. Equation (4.31)). Differentiating with respect to ϕ yields

$$\frac{d\mathbf{K}}{d\phi}\mathbf{d} + \mathbf{K}\frac{d\mathbf{d}}{d\phi} = \frac{d\mathbf{f}}{d\phi}, \quad \Rightarrow \quad \frac{d\mathbf{d}}{d\phi} = \mathbf{K}^{-1} \left(\frac{d\mathbf{f}}{d\phi} - \frac{d\mathbf{K}}{d\phi}\mathbf{d} \right). \quad (4.36)$$

The compliance is $c(\phi, \mathbf{d}) = \mathbf{f}^T \mathbf{d}$. Therefore,

$$\frac{dc_{\text{eq}}}{d\phi} = \left(\frac{d\mathbf{f}}{d\phi} \right)^T \mathbf{d} + \mathbf{f}^T \frac{d\mathbf{d}}{d\phi}. \quad (4.37)$$

Using $\mathbf{f} = \mathbf{K}\mathbf{d}$ and $\mathbf{K} = \mathbf{K}^T$ (self-adjoint case), one obtains

$$\frac{dc_{\text{eq}}}{d\phi} = 2 \mathbf{d}^T \frac{d\mathbf{f}}{d\phi} - \mathbf{d}^T \frac{d\mathbf{K}}{d\phi} \mathbf{d}. \quad (4.38)$$

In the standard minimum-compliance setting, the external load vector \mathbf{f} is assumed independent of ϕ , so that $d\mathbf{f}/d\phi = \mathbf{0}$ and

$$\frac{dc_{\text{eq}}}{d\phi} = -\mathbf{d}^T \frac{d\mathbf{K}}{d\phi} \mathbf{d}. \quad (4.39)$$

The stationarity condition with respect to λ enforces the volume constraint:

$$\frac{\partial \mathcal{L}_{\text{eq}}}{\partial \lambda} = v(\phi) = 0, \quad \text{i.e.,} \quad \sum_{e=1}^{n_e} \phi_e V_e - fV_0 = 0. \quad (4.40)$$

The discrete problem is solved iteratively by: (i) solving $\mathbf{K}(\phi)\mathbf{d} = \mathbf{f}$, (ii) updating ϕ according to the optimality condition, and (iii) updating λ (e.g., via bisection) to satisfy the volume constraint.

4.2.2.4 Optimality Criterion

In the discrete formulation of topology optimization, the optimality criterion is used to iteratively update the material density distribution ϕ across the design domain. The goal is to adjust ϕ in a manner that progressively minimizes compliance while satisfying the volume constraint.

To accomplish this, an update rule for ϕ is established at each iteration, using a trial value, denoted as ϕ_{trial} , which incorporates information on the local optimality condition.

The trial value of the design variable ϕ_{trial} at iteration $k + 1$ is computed as:

$$\phi_{\text{trial}}^{k+1} = \phi^k [B_{\text{eq}}(\phi^k, \lambda^{k+1})]^\eta, \quad (4.41)$$

where ϕ^k represents the material density at iteration k and $B_{\text{eq}}(\phi^k, \lambda^{k+1})$ is the control variable that measures the local sensitivity of compliance with respect to ϕ :

$$B_{\text{eq}}(\phi^k, \lambda^{k+1}) = -\frac{\partial c_{\text{eq}}(\phi^k)}{\partial \phi^k} \left(\lambda^{k+1} \frac{\partial v(\phi)}{\partial \phi^k} \right)^{-1}, \quad (4.42)$$

where λ^{k+1} is the Lagrange multiplier associated with the volume constraint at iteration $k + 1$ and η is a damping factor, typically chosen as 0.5 to ensure numerical stability [231, 251].

The control variable B_{eq} approaches unity when the compliance sensitivity balances with the volume constraint sensitivity, satisfying the optimality condition locally. Therefore, when $B_{\text{eq}} = 1$, the material density does not change. If $B_{\text{eq}} \neq 1$, ϕ_{trial} is either increased or decreased accordingly, ensuring a smooth and directed adjustment of the material distribution towards optimality.

To control the update step and prevent excessive changes in material density that could destabilize the optimization process, an auxiliary variable ϕ_{aux} is introduced. This variable restricts the update step size between iterations, enhancing stability. The auxiliary variable ϕ_{aux} is defined as:

$$\phi_{\text{aux}} = \begin{cases} \phi^k - m, & \text{if } \phi_{\text{trial}} \leq \phi^k - m, \\ \phi^k + m, & \text{if } \phi_{\text{trial}} \geq \phi^k + m, \\ \phi_{\text{trial}}, & \text{otherwise,} \end{cases} \quad (4.43)$$

where m is the move limit, a positive constant that restricts the maximum allowable change in ϕ (often $m = 0.2$ in applications [231, 251]). This constraint on step size prevents abrupt variations in the material density, facilitating a stable progression towards the solution.

Once the auxiliary variable ϕ_{aux} is computed, the material density is further projected to ensure that it remains within the physical bounds $[0, 1]$. The projection is performed as:

$$\phi_{\text{project}} = \min(\max(\phi_{\text{aux}}, 0), 1), \quad (4.44)$$

where ϕ_{project} is constrained within $[0, 1]$, thus ensuring the material density is either void ($\phi = 0$) or fully solid ($\phi = 1$).

To improve convergence and ensure a manufacturable design, a filtering step is applied to the updated material density distribution. Filtering, as discussed in Section 4.2.2.5.

In summary, the optimality criterion iteratively updates the material densities in a multi-step process:

1. Compute the trial value ϕ_{trial} using the equilibrium control variable B_{eq} .
2. Limit the changes in ϕ using the auxiliary variable ϕ_{aux} to ensure stability.
3. Project ϕ_{aux} into the physical bounds $[0, 1]$ to maintain feasible values.
4. Apply a filter to ensure smoothness and manufacturability.

This iterative process continues until the material distribution ϕ converges, yielding a configuration that minimizes compliance while adhering to both equilibrium and volume constraints.

4.2.2.5 Regularization Techniques

In the discrete formulation of topology optimization, regularization techniques [32, 220, 245–248] are essential to ensure that the optimization process yields stable, physically meaningful, and manufacturable solutions. Without proper regularization, the optimization may produce numerical artefacts such as checkerboarding, mesh dependency, and highly localized material concentrations, which detract from the quality and practicality of the design [249].

As discussed in Section 4.2.1.5, Radial Basis Function (RBF) filtering [246] is one of the most effective techniques for smoothing the distribution of both material densities and sensitivities across neighbouring elements. This approach reduces abrupt transitions between solid and void regions, ensuring a stable optimization process that converges to a manufacturable design.

The filtered density for each element e at iteration $k + 1$ is computed as a weighted average of the densities of neighbouring elements:

$$\phi_e^{k+1} = \frac{\sum_{j \in \mathcal{N}(e)} \kappa(r(\mathbf{x}_e, \mathbf{x}_j)) \phi_{j, \text{project}}}{\sum_{j \in \mathcal{N}(e)} \kappa(r(\mathbf{x}_e, \mathbf{x}_j))}, \quad (4.45)$$

where $\mathcal{N}(e)$ denotes the set of neighbouring elements surrounding element e , $\kappa(r(\mathbf{x}_e, \mathbf{x}_j))$ is the RBF filter function, and $r(\mathbf{x}_e, \mathbf{x}_j)$ represents the normalized distance between the centroids of elements e and j :

$$r(\mathbf{x}_e, \mathbf{x}_j) = \frac{\|\mathbf{x}_e - \mathbf{x}_j\|}{r_{\text{min}}}, \quad (4.46)$$

where r_{min} is the filter radius that controls the influence range of the filter function, ensuring that only elements within a specified distance contribute to the filtered density of a given element.

The compactly supported RBF filter [246] is defined as:

$$\kappa(r(\mathbf{x}_e, \mathbf{x}_j)) = [r_0(\mathbf{x}_e, \mathbf{x}_j)]^4 (4r(\mathbf{x}_e, \mathbf{x}_j) + 1), \quad (4.47)$$

where $r_0(\mathbf{x}_e, \mathbf{x}_j) = \max(0, 1 - r(\mathbf{x}_e, \mathbf{x}_j))$ ensures that only elements within the specified radius r_{min} affect the filtered density, resulting in localized smoothing.

The density filtering process is crucial for producing smooth material distributions and preventing the formation of non-physical patterns, such as checkerboarding, across the finite element mesh.

In alternative to density filtering, sensitivity filtering is often applied to smooth the sensitivities of compliance with respect to the material densities. This filtering step stabilizes the optimization by ensuring that the sensitivity values vary smoothly across the design domain, preventing localized irregularities in the density distribution.

The filtered sensitivity $\frac{\partial c_{\text{eq}}}{\partial \phi_e}$ for each element e is computed analogously to density filtering:

$$\frac{\partial c_{\text{eq}}}{\partial \phi_e} = \frac{\sum_{j \in \mathcal{N}(e)} \kappa(r(\mathbf{x}_e, \mathbf{x}_j)) \frac{\partial c_{\text{eq}}}{\partial \phi_{j, \text{project}}}}{\sum_{j \in \mathcal{N}(e)} \kappa(r(\mathbf{x}_e, \mathbf{x}_j))}. \quad (4.48)$$

Sensitivity filtering prevents abrupt changes in the compliance sensitivities, which could destabilize the optimization process or lead to convergence on non-optimal designs. Additionally, filtering the sensitivities enhances the robustness of the optimization, particularly when dealing with high-resolution meshes or complex design domains.

In conclusion, the advantages of regularization techniques in the discrete setting are numerous:

- **Prevention of Checkerboarding.** Filtering smooths the material distribution, eliminating non-physical checkerboard patterns that can emerge in element-based optimization.
- **Reduction of Mesh Dependency.** Regularization reduces the dependency of the optimization results on mesh size or element shape, yielding designs that are more robust and consistent across different mesh discretizations.
- **Enhanced Manufacturability.** The smoothed material distribution resulting from filtering is more likely to produce designs feasible for manufacturing, as it avoids sharp, localized transitions between solid and void regions.
- **Stabilization of the Optimization Process.** By smoothing the sensitivities, filtering helps the optimization process to progress in a stable manner, avoiding sudden, destabilizing changes in material distribution.

In summary, regularization techniques are an indispensable part of the discrete topology optimization process. It ensures that the final design is physically meaningful, manufacturable, and free from numerical artefacts, thus enhancing the reliability and applicability of the optimized structures.

4.2.3 Numerical Implementation

In the present section, we describe the algorithm implemented to solve the topology optimization problem using the Solid Isotropic Material with Penalization (SIMP) method. The formulation follows the standard SIMP-based topology optimization framework combined with an Optimality Criterion (OC) update scheme, as widely established in the literature [32, 220].

The implementation is carried out using the AceGen and AceFEM packages in Mathematica[®]. Specifically, AceFEM is used to manage the optimization process and finite element analysis, while AceGen is employed to generate the finite elements required for the discretization of the design domain.

The topology optimization algorithm implemented in AceFEM consists of six main stages:

- *General Parameters*. This stage defines the problem's fundamental parameters, including geometry, material properties, loading conditions, SIMP method parameters (e.g., penalization factor, filter radius), and convergence criteria.
- *Finite Element Analysis Initialization*. In this step, the design domain geometry, material properties, meshing strategy (using finite elements generated by AceGen), and boundary conditions are specified.
- *SIMP Analysis Initialization*. Here, all essential quantities for the SIMP method are initialized, including the design variables (material density), neighborhood of elements (for filtering), and initial compliance values.
- *Density Filter Computation*. The density filter is computed at this stage using Radial Basis Functions (RBF), as outlined in Section 4.2.2.5. This filter smooths the density field, promoting numerical stability.
- *Topology Optimization Algorithm*. This phase corresponds to the implementation of the SIMP-based optimization loop combined with the OC update strategy, as detailed in Algorithm 1.
- *Solution and Visualization of Results*. Finally, the optimized material distribution is computed, and the results are visualized.

As discussed in Section 4.2.2.5, both density-based and sensitivity-based filtering techniques can be used to regularize the solution of the topology optimization problem. Since these methods differ only in the variable to which the filter is applied, the algorithm below presents the density-based filtering technique as an illustrative example.

The following algorithmic framework describes the implementation of a standard SIMP-based topology optimization procedure. Algorithm 1 presents the main iterative loop, where the equilibrium displacements are solved and material densities are updated. Algorithm 2 enforces the volume constraint using a bisection method to determine the Lagrange multiplier λ , while Algorithm 3 updates the design variable ϕ using the optimality criterion combined with projection and filtering steps.

This structured approach provides a consistent implementation of the SIMP-based topology optimization procedure, enabling the computation of material distributions that satisfy equilibrium and volume constraints.

Algorithm 1: Topology Optimization Solution

- 1 **Input:** $\lambda_{\text{old}} = \lambda^k, \phi_{\text{old}} = \phi^k, \text{maxIter}$
 - 2 **Output:** $\lambda_{\text{new}} = \lambda^{k+1}, \phi_{\text{new}} = \phi^{k+1}$
 - 3 **for** $k = 0$ **to** maxIter **do**
 - 4 Solve $\mathbf{Kd} = \mathbf{f}$ to compute the equilibrium displacement field \mathbf{d} for the current material distribution.
 - 5 Apply Algorithm 2 to compute the updated values of λ^{k+1} and ϕ^{k+1} based on the optimality criterion and volume constraint.
 - 6 Update the values: $\phi^k \leftarrow \phi^{k+1}, \lambda^k \leftarrow \lambda^{k+1}$ for the next iteration.
-

Algorithm 2: Bisection Method for Volume Constraint Enforcement

- 1 **Input:** $\lambda^k, \phi^k, f, V, \lambda_a, \lambda_b, \text{tol}$
 - 2 **Output:** $\lambda^{k+1}, \phi^{k+1}$
 - 3 Compute $c_{\text{eq}}(\lambda^k, \phi^k)$ and its sensitivity $\frac{\partial c_{\text{eq}}(\lambda^k, \phi^k)}{\partial \phi}$
 - 4 Compute the volume $v(\phi^k)$ and its sensitivity $\frac{\partial v(\phi^k)}{\partial \phi}$
 - 5 Execute **Algorithm 3** to update ϕ^{k+1} and $v(\phi^{k+1})$
 - 6 **while** $\lambda_b - \lambda_a > \text{tol}$ **do**
 - 7 Set $\lambda^{k+1} = \frac{\lambda_a + \lambda_b}{2}$
 - 8 Re-evaluate **Algorithm 3** to compute ϕ^{k+1} and $v(\phi^{k+1})$
 - 9 **if** $v(\phi^{k+1}) - fV = 0$ **then**
 - 10 **break**
 - 11 **else**
 - 12 **if** $v(\phi^{k+1}) - fV > 0$ **then**
 - 13 Set $\lambda_a = \lambda^{k+1}$
 - 14 **else**
 - 15 Set $\lambda_b = \lambda^{k+1}$
-

4.2.4 Numerical Examples

In present section, we present several numerical examples to demonstrate the application of the SIMP-based topology optimization method as formulated in Sections 4.2.1 and 4.2.2. Each example applies density-based and sensitivity-based filtering techniques, as discussed in Sections 4.2.1.5 and 4.2.2.5, to ensure stability and manufacturability in the resulting designs.

These examples illustrate the versatility and efficacy of the SIMP method across different

Algorithm 3: Design Variable Update using Optimality Criterion

- 1 **Input:** $\phi^k, c_{\text{eq}}, f, V, m, \eta, \mathcal{K}, \lambda^{k+1}$
 - 2 **Output:** ϕ^{k+1}, v^{k+1}
 - 3 **if** $\lambda^{k+1} > 0$ **then**
 - 4 Calculate the equilibrium control variable $B_{\text{eq}}(\phi^k, \lambda^{k+1})$
 - 5 Set $\phi_{\text{trial}} = \phi^k B_{\text{eq}}(\phi^k, \lambda^{k+1})\eta$
 - 6 Compute the auxiliary variable ϕ_{aux} to limit the step size
 - 7 Project ϕ_{aux} into the range $[0, 1]$ to obtain ϕ_{proj}
 - 8 Apply the density filter \mathcal{K} to ϕ_{proj} to smooth the design variable, yielding ϕ^{k+1}
 - 9 Compute the total volume $v(\phi^{k+1})$
-

structural configurations, loading conditions, and dimensionalities, drawing on benchmark problems from the literature [3, 251]. The material properties for all examples assume isotropic A-36 steel, characterized by a Young's modulus of 2.11×10^{11} Pa, a Poisson's ratio of 0.29, and a yield stress $\sigma_y = 2.2 \times 10^8$ Pa.

Table 4.1 summarizes the main parameters used in the topology optimization process, which include the penalization factor p , damping factor η , move limit m , volume fraction f , filter radius r_{min} , maximum iterations, and tolerance levels.

Table 4.1: Parameters used in the topology optimization process.

Parameter	Symbol	Value
Penalization factor	p	3
Damping factor	η	0.5
Move limit	m	0.2
Volume fraction	f	0.4
Filter radius	r_{min}	0.05
Maximum iterations	<code>maxIter</code>	150
Tolerance	<code>tol</code>	1×10^{-6}
λ lower/upper limit	λ_a/λ_b	0/10 ⁵

Each example is presented with a detailed description of the geometry, boundary conditions, applied loads, and mesh discretization. The results are analyzed to evaluate the structural performance and material efficiency of each optimized design.

4.2.4.1 Cantilever Beam

The first example involves a classic 2D cantilever beam (Figure 4.1) with dimensions of 2 m in length and 1 m in height. The left end of the beam is fixed, and different loading conditions are applied at the right end. Three load configurations are considered: (a) a central load, (b) a load applied at the bottom-right corner, and (c) a load applied at the top-right corner. Table 4.2 lists the load magnitudes for each configuration.

The design domain is discretized using a mesh of $160 \times 80 = 12.800$ square elements, each with an area of $0.0125 \times 0.0125 \text{ m}^2$, providing adequate resolution to capture the response under each loading condition.

Table 4.2: Load values for the 2D cantilever beam under different loading conditions [3].

Problem	Load Position	Load Magnitude (P)
Cantilever Beam	Central load	$8.6 \times 10^4 \text{ N}$
	Bottom-right corner	$6.2 \times 10^4 \text{ N}$
	Top-right corner	$6.2 \times 10^4 \text{ N}$

4.2.4.2 2D Short Cantilever Beam

The second example studies a short cantilever beam with dimensions of 1 m in length and 1 m in height. The left side is fixed, and loads are applied at various locations on the right end. Three loading scenarios are considered: (a) central load, (b) load at bottom-right corner, and (c) load at top-right corner (Figure 4.2). The load magnitudes are listed in Table 4.3.


The design domain is discretized into $80 \times 80 = 6.400$ square elements, each with an area of $0.0125 \times 0.0125 \text{ m}^2$.

Table 4.3: Load magnitudes for the short cantilever beam under various loading configurations [3].


Problem	Load Position	Load Magnitude (P)
Short Cantilever Beam	Central load	$9.0 \times 10^4 \text{ N}$
	Bottom-right corner	$5.8 \times 10^4 \text{ N}$
	Top-right corner	$5.8 \times 10^4 \text{ N}$

4.2.4.3 L-Shaped Structure


The third example examines a 2D L-shaped structure (Figure 4.3), with fixed boundaries and applied loads. Two loading conditions are considered: (a) a central load and (b) a load at the bottom-right corner. The load magnitudes are given in Table 4.4.



(a)




(b)




(c)

Figure 4.1: 2D cantilever beam with various load applications [3]. (a) Central load. (b) Load at bottom-right corner. (c) Load at top-right corner.


The design domain is discretized into 14,400 square elements, each measuring $0.0125 \times$



(a)




(b)




(c)

Figure 4.2: Short cantilever beam with different loading configurations [3]. (a) Central load. (b) Load at bottom-right corner. (c) Load at top-right corner.

0.0125 m², ensuring a detailed structural response.



(a)



(b)

Figure 4.3: L-shaped structure under different loading configurations [3]. (a) Central load. (b) Load at bottom-right corner.

Table 4.4: Load magnitudes for the L-shaped structure under different loading conditions [3].

Problem	Load Position	Load Magnitude (P)
L-Shaped Structure	Central load	1.5×10^4 N
	Bottom-right corner	1.5×10^4 N

4.2.5 Results and Discussion

In the present section, the results obtained from the topology optimization problems introduced in Section 4.2.4 are analysed. For each benchmark problem, the performance of density-

based filtering (DBF) and sensitivity-based filtering (SBF) techniques is evaluated in terms of convergence behaviour, compliance values, and resulting material distributions. The results are compared with benchmark data from the literature [3] in order to assess the accuracy and consistency of the implementation.

4.2.5.1 Cantilever Beam

The optimized material distributions for the 2D cantilever beam under three loading scenarios—central load, bottom-right corner load, and top-right corner load—are shown in Figure 4.4. The resulting topologies exhibit well-defined load paths, with material distributed along the principal stress trajectories. These results are consistent with classical solutions reported in the literature, indicating that the SIMP-based formulation accurately captures the structural response under different loading configurations.

The quantitative results reported in Table 4.5 show that both DBF and SBF approaches produce compliance values in close agreement with the benchmark data. The number of iterations required for convergence is slightly lower for DBF in all cases, indicating a modest computational advantage while maintaining comparable accuracy.


Table 4.5: Comparison of DBF and SBF filtering strategies for the cantilever beam. The compliance values closely match benchmark results, while DBF shows a slightly faster convergence in terms of iteration count.

Test Problem	Data	DBF	SBF	Benchmark
Central load	Iterations	40	43	42
	Compliance	2.3087×10^2	2.3087×10^2	2.3086×10^2
Bottom load	Iterations	38	40	39
	Compliance	1.2912×10^2	1.2912×10^2	1.2910×10^2
Top load	Iterations	38	39	39
	Compliance	1.2910×10^2	1.2910×10^2	1.2910×10^2


4.2.5.2 2D Short Cantilever Beam

The optimized topologies for the short cantilever beam under the three loading configurations are reported in Figure 4.5. The material distributions follow the expected structural patterns, with material concentrated along the dominant load paths. Despite the reduced domain size, the method captures the essential features of the optimal layout.


The numerical results in Table 4.6 confirm a strong agreement with the benchmark data. Both filtering techniques lead to nearly identical compliance values, while DBF again requires fewer iterations in most cases. These results indicate that both approaches provide stable and accurate solutions.



(a) Central load configuration



(b) Bottom-right corner load configuration




(c) Top-right corner load configuration


Figure 4.4: Optimized material distributions for the cantilever beam under different loading configurations. The layouts obtained using the SIMP method highlight the formation of efficient load-carrying paths, with material concentrated along principal stress directions for each loading case.

4.2.5.3 L-Shaped Structure


The results for the L-shaped structure under two loading conditions are shown in Figure 4.6. The optimized layouts exhibit clear structural patterns, with material distributed along load-bearing



(a) Central load configuration



(b) Bottom-right corner load configuration



(c) Top-right corner load configuration

Figure 4.5: Optimized material distributions for the short cantilever beam under different loading configurations. The results demonstrate consistent identification of efficient load paths, even in a reduced design domain.

paths and around stress concentration regions. The method effectively handles the geometric

Table 4.6: Comparison of DBF and SBF filtering strategies for the short cantilever beam. The results show excellent agreement with benchmark values and consistent convergence behaviour.

Test Problem	Data	DBF	SBF	Benchmark
Central load	Iterations	25	27	27
	Compliance	5.6124×10^1	5.6124×10^1	5.6125×10^1
Bottom load	Iterations	21	23	23
	Compliance	2.8287×10^1	2.8287×10^1	2.8285×10^1
Top load	Iterations	21	23	22
	Compliance	2.8287×10^1	2.8287×10^1	2.8286×10^1

complexity of the domain.

The quantitative comparison in Table 4.7 shows that both DBF and SBF approaches produce compliance values that closely match the benchmark data. As observed in the previous examples, DBF generally converges in fewer iterations, while both methods achieve comparable solution accuracy.


Table 4.7: Comparison of DBF and SBF filtering strategies for the L-shaped structure. The results confirm the consistency of both approaches with benchmark data, with minor differences in convergence behaviour.

Test Problem	Data	DBF	SBF	Benchmark
Central load	Iterations	27	29	30
	Compliance	2.907×10^1	2.907×10^1	2.906×10^1
Bottom load	Iterations	25	25	29
	Compliance	2.0193×10^1	2.0193×10^1	2.0191×10^1


4.3 A Comprehensive Theoretical and Algorithmic Study on the Method of Moving Asymptotes for Structural Optimization

The present section provides a comprehensive theoretical and algorithmic investigation of the MMA for structural optimization problems. The focus is placed on the analytical foundations of the method, its algorithmic structure, and its numerical behaviour when applied to constrained and large-scale optimization problems arising in structural mechanics.

Section 4.3.1 formulates the general constrained structural optimization problem. Section 4.3.2 presents the theoretical foundations of the MMA. Section 4.3.3 details the algorithmic framework



(a) Central load configuration



(b) Bottom-right corner load configuration

Figure 4.6: Optimized material distributions for the L-shaped structure under different loading configurations. The results highlight the ability of the method to capture complex load transfer mechanisms in non-rectangular domains.

and implementation aspects. Finally, Section 4.3.4 reports the numerical validation through benchmark tests.

4.3.1 Constrained Structural Optimization Problem: Formulation and Analytical Motivation

The present section formulates the general constrained structural optimization problem and analyses the main analytical challenges associated with nonlinear and nonconvex problems. The discussion highlights how the analytical challenges motivate the introduction of local and separable approximation strategies that form the analytical foundation of the MMA. Section 4.3.1.1 presents the general mathematical formulation of constrained structural optimization, defining

the objective function, design variables, and equality and inequality constraints. Section 4.3.1.2 examines the analytical challenges in nonlinear and nonconvex problems. Finally, Section 4.3.1.3 discusses the motivation for adopting local and separable convex approximations, introducing the conceptual bridge toward the theoretical formulation of the MMA.

4.3.1.1 General Mathematical Formulation

A general constrained structural optimization problem can be formulated as:

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}_{\text{adm}}^A} \quad & f(\mathbf{x}) \\ \text{subject to} \quad & \mathbf{h}(\mathbf{x}) = \mathbf{0} \\ & \mathbf{g}(\mathbf{x}) \leq \mathbf{0} \end{aligned} \tag{4.49}$$

where $\mathbf{x} = [x_1, x_2, \dots, x_A]^T \in \mathbb{R}^A$ denotes the vector of design variables, and $f : \mathbb{R}^A \rightarrow \mathbb{R}$ is the objective function to be minimized. The functions $\mathbf{h} : \mathbb{R}^A \rightarrow \mathbb{R}^B$ and $\mathbf{g} : \mathbb{R}^A \rightarrow \mathbb{R}^C$ collect the equality and inequality constraints, respectively, which enforce equilibrium equations, boundary conditions, and additional design requirements. The admissible design domain is defined by the box constraints:

$$\mathbb{R}_{\text{adm}}^A = \{ \mathbf{x} \in \mathbb{R}^A \mid \mathbf{x}^{\min} \leq \mathbf{x} \leq \mathbf{x}^{\max} \} \tag{4.50}$$

where $\mathbf{x}^{\min}, \mathbf{x}^{\max} \in \mathbb{R}^A$ represent the lower and upper bounds on the design variables.

The objective function $f(\mathbf{x})$ quantifies the structural performance to be optimized — such as compliance, stiffness, weight, or stress — depending on the specific problem formulation. The constraint functions $\mathbf{h}(\mathbf{x})$ and $\mathbf{g}(\mathbf{x})$ define the feasible design space by enforcing mechanical admissibility and physical consistency. The dimensionality A of the design space can range from a few to several thousand variables, particularly in topology optimization, where each finite element or material density fraction represents a design variable [218, 234, 252–256].

The constrained structural optimization problem in Equation (4.49) is, in general, nonlinear and nonconvex. The analytical difficulties relate to the nonlinearity and nonconvexity of the optimization problem in Equation (4.49) are discussed in Section 4.3.1.2.

4.3.1.2 Analytical Difficulties in Nonlinear and Nonconvex Problems

Nonlinearity and nonconvexity constitute the main analytical difficulties in the constrained structural optimization problem defined in Equation (4.49) [257, 258]. Both phenomena modify the mathematical structure of the problem by affecting the regularity and curvature of the objective and constraint functions within the admissible domain. Their combined effect complicates the analytical characterization of the solution and prevents the direct establishment of conditions guaranteeing existence, uniqueness, and stability.

The problem in Equation (4.49) is said to be *nonlinear* when the objective function $f(\mathbf{x})$ depends nonlinearly on the design variables $\mathbf{x} \in \mathbb{R}_{\text{adm}}^A$ (Equation (4.50)), as a result of geometric

nonlinearities or nonlinear material constitutive behaviour [259]. Such dependence affects the local behaviour of the first- and second-order derivatives, defined as

$$\nabla f(\mathbf{x}) = \left[\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, \dots, \frac{\partial f}{\partial x_A} \right]^\top \quad \nabla^2 f(\mathbf{x}) = \left[\frac{\partial^2 f}{\partial x_i \partial x_j} \right]_{i,j=1}^A \quad (4.51)$$

Analogous definitions apply to the constraint functions $\mathbf{h}(\mathbf{x})$ and $\mathbf{g}(\mathbf{x})$, whose first- and second-order derivatives (gradients and Hessians) govern the curvature and differentiability properties of the feasible domain. Strong nonlinearities may induce rapid spatial variations in the derivatives or abrupt curvature changes in the Hessians, leading to irregular behaviour and possible loss of smoothness within the admissible domain.

In such cases, the assumptions of differentiability and Lipschitz continuity — fundamental to regularity analysis in optimization — may no longer hold, thus reducing the analytical tractability of the problem [260].

Nonconvexity introduces an additional and more fundamental source of difficulty. Let $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}_{\text{adm}}^A$ denote two admissible configurations of the design variables. The objective function $f(\mathbf{x})$ is convex if it satisfies

$$f(\mathbf{x}_2) \geq f(\mathbf{x}_1) + \nabla f(\mathbf{x}_1)^\top (\mathbf{x}_2 - \mathbf{x}_1) \quad \forall \mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}_{\text{adm}}^A \quad (4.52)$$

When Equation (4.52) does not hold, the function is nonconvex and its curvature cannot be bounded by a positive semidefinite Hessian, i.e.,

$$\nabla^2 f(\mathbf{x}) < 0 \quad \text{for some } \mathbf{x} \in \mathbb{R}_{\text{adm}}^A \quad (4.53)$$

A similar reasoning applies to the constraint functions $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$, whose loss of convexity alters the geometry of the feasible domain, possibly resulting in multiple disconnected feasible regions. As a consequence, the admissible domain may contain several stationary points $\nabla f(\mathbf{x}^*) = \mathbf{0}$, corresponding to local minima, maxima, or saddle points. Geometrically, the curvature of $f(\mathbf{x})$, $\mathbf{g}(\mathbf{x})$, and $\mathbf{h}(\mathbf{x})$ may vary sharply across the domain, producing irregular configurations where analytical consistency is lost.

From a theoretical standpoint, nonlinearity primarily affects the differentiability and conditioning of the functions involved, whereas nonconvexity alters their global curvature and compromises the stability of stationary solutions.

4.3.1.3 Analytical Motivation for Local and Separable Approximations

The analytical difficulties discussed in Section 4.3.1.2 motivate the reformulation of the constrained structural optimization problem in Equation (4.49) into auxiliary problems exhibiting well-defined convexity properties and enhanced analytical tractability. A common approach consists in constructing, at a generic reference configuration $\bar{\mathbf{x}} \in \mathbb{R}_{\text{adm}}^A$, locally convex and separable approximations of the objective function $f(\mathbf{x})$ and of the constraint functions $\mathbf{h}(\mathbf{x})$ and $\mathbf{g}(\mathbf{x})$. These local approximations retain the first-order information of the original functions while introducing curvature control through explicit curvature parameters.

Given the reference configuration $\bar{\mathbf{x}}$, the local approximations of the objective function $\tilde{f}(\mathbf{x}; \bar{\mathbf{x}})$ and of the constraint functions $\tilde{\mathbf{g}}(\mathbf{x}; \bar{\mathbf{x}})$ and $\tilde{\mathbf{h}}(\mathbf{x}; \bar{\mathbf{x}})$ are constructed to satisfy first-order consistency conditions:

$$\begin{aligned}\tilde{f}(\bar{\mathbf{x}}; \bar{\mathbf{x}}) &= f(\bar{\mathbf{x}}) & \nabla_{\mathbf{x}} \tilde{f}(\bar{\mathbf{x}}; \bar{\mathbf{x}}) &= \nabla f(\bar{\mathbf{x}}) \\ \tilde{\mathbf{g}}(\bar{\mathbf{x}}; \bar{\mathbf{x}}) &= \mathbf{g}(\bar{\mathbf{x}}) & \nabla_{\mathbf{x}} \tilde{\mathbf{g}}(\bar{\mathbf{x}}; \bar{\mathbf{x}}) &= \nabla \mathbf{g}(\bar{\mathbf{x}}) \\ \tilde{\mathbf{h}}(\bar{\mathbf{x}}; \bar{\mathbf{x}}) &= \mathbf{h}(\bar{\mathbf{x}}) & \nabla_{\mathbf{x}} \tilde{\mathbf{h}}(\bar{\mathbf{x}}; \bar{\mathbf{x}}) &= \nabla \mathbf{h}(\bar{\mathbf{x}})\end{aligned}\quad (4.54)$$

ensuring that the approximations reproduce both the function values and gradients of the original problem in the neighbourhood of $\bar{\mathbf{x}}$. Equation (4.54) guarantees first-order consistency and preserves the differential information required to approximate the stationary behaviour of the constrained problem in Equation (4.49).

To simplify the analytical dependence on the design variables $\mathbf{x} = [x_1, x_2, \dots, x_A]^\top$, the local approximations adopt a separable structure:

$$\begin{aligned}\tilde{f}(\mathbf{x}; \bar{\mathbf{x}}) &= \rho_f(\bar{\mathbf{x}}) + \sum_{i=1}^A \varphi_i(x_i; \bar{\mathbf{x}}) \\ \tilde{\mathbf{g}}(\mathbf{x}; \bar{\mathbf{x}}) &= \rho_{\mathbf{g}}(\bar{\mathbf{x}}) + \sum_{i=1}^A \psi_i(x_i; \bar{\mathbf{x}}) \\ \tilde{\mathbf{h}}(\mathbf{x}; \bar{\mathbf{x}}) &= \rho_{\mathbf{h}}(\bar{\mathbf{x}}) + \sum_{i=1}^A \chi_i(x_i; \bar{\mathbf{x}})\end{aligned}\quad (4.55)$$

where each term φ_i , ψ_i , and χ_i depends on a single design variable x_i . The coefficients ρ_f , $\rho_{\mathbf{g}}$, and $\rho_{\mathbf{h}}$ are determined so as to satisfy the conditions in Equation (4.54). The separable representation in Equation (4.55) removes the coupling among design variables and yields diagonal Hessian matrices for the local approximations, thereby simplifying the analysis of curvature and the verification of convexity.

Convexity is enforced by prescribing explicit curvature bounds for each univariate component. For the objective function, strict convexity is ensured if

$$\frac{d^2}{dx_i^2} \varphi_i(x_i; \bar{\mathbf{x}}) \geq m_{f,i}(\bar{\mathbf{x}}) > 0 \quad x_i \in [x_i^{\min}, x_i^{\max}] \quad i = 1, \dots, A \quad (4.56)$$

which guarantees $\nabla_{\mathbf{x}}^2 \tilde{f}(\mathbf{x}; \bar{\mathbf{x}}) \geq \text{diag}(m_{f,1}, \dots, m_{f,A}) > \mathbf{0}$. Analogous curvature conditions apply to the constraint functions:

$$\frac{d^2}{dx_i^2} \psi_i(x_i; \bar{\mathbf{x}}) \geq \mathbf{m}_{\mathbf{g},i}(\bar{\mathbf{x}}) \geq \mathbf{0} \quad \frac{d^2}{dx_i^2} \chi_i(x_i; \bar{\mathbf{x}}) \geq \mathbf{m}_{\mathbf{h},i}(\bar{\mathbf{x}}) \geq \mathbf{0} \quad (4.57)$$

The vector parameters $\mathbf{m}_{\mathbf{g},i}$ and $\mathbf{m}_{\mathbf{h},i}$ control the curvature of the local constraint approximations. Equations (4.56) and (4.57) introduce an explicit analytical mechanism for curvature control,

directly addressing the irregular behaviour described in Section 4.3.1.2 and ensuring local convexity in compliance with Equation (4.52).

The local convex and separable approximations defined by Equation (4.55) lead to the subproblem associated with the reference configuration $\bar{\mathbf{x}}$:

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}_{\text{adm}}^A} \quad & \tilde{f}(\mathbf{x}; \bar{\mathbf{x}}) \\ \text{subject to} \quad & \tilde{\mathbf{h}}(\mathbf{x}; \bar{\mathbf{x}}) = \mathbf{0} \\ & \tilde{\mathbf{g}}(\mathbf{x}; \bar{\mathbf{x}}) \leq \mathbf{0} \end{aligned} \tag{4.58}$$

where the admissible domain is defined in Equation (4.50). Under the first-order consistency conditions (Equation (4.54)) and curvature control constraints (Equations (4.56) and (4.57)), the constrained problem in Equation (4.58) admits a unique minimizer within the admissible domain. The separable and convex formulation ensures mathematical stability, analytical clarity, and well-posed curvature behaviour of the local approximation.

The construction of local convex and separable subproblems forms the analytical foundation of sequential explicit convex approximation methods [261, 262]. Within this class, the MMA [4, 33, 237] provides a general and effective formulation.

4.3.2 The Method of Moving Asymptotes: Theoretical Foundations

The section presents the theoretical foundations of the MMA and examines its analytical structure in detail.

Section 4.3.2.1 introduces the MMA formulation for a one-dimensional unconstrained structural optimization problem, where the main analytical concepts can be derived and interpreted mathematically. Section 4.3.2.2 extends the formulation to a one-dimensional constrained structural optimization problem. Finally, Section 4.3.2.3 generalizes the method to multidimensional and large-scale structural optimization problem.

4.3.2.1 One-Dimensional Unconstrained Structural Optimization Problem

To isolate the basic foundations of the MMA, the following assumptions are introduced into the constrained structural optimization problem in Equation (4.49):

1. the design space is one-dimensional, with a single scalar variable ($A = 1$ in Equation (4.50)), which permits a fully analytical treatment of the optimization problem;
2. no equality constraints are imposed, i.e., relations of the form $\mathbf{h}(\mathbf{x}) = \mathbf{0}$ are absent;
3. no inequality constraints are imposed, i.e., conditions of the form $\mathbf{g}(\mathbf{x}) \leq \mathbf{0}$ are not considered;
4. the objective function is assumed to be nonlinear and nonconvex within the admissible design domain (see Sections 4.3.1.2 and 4.3.1.3).

Under the indicated assumptions, the constrained structural optimization problem in Equation (4.49) reduces to the nonlinear and nonconvex unconstrained optimization problem:

$$\min_{x \in (x^{\min}, x^{\max})} f(x) \quad (4.59)$$

where $x^{\min}, x^{\max} \in \mathbb{R}$ denote the lower and upper bounds on the admissible design domain (Equation (4.50)).

4.3.2.1.1 Construction of the Convex Approximation

Let $x \in \mathbb{R}$ denote the scalar design variable introduced in the one-dimensional unconstrained optimization problem (see Equation (4.59)). The local convex approximation of the nonlinear and nonconvex objective function $f(x)$ is constructed within an open interval $(L, U) \subset (x^{\min}, x^{\max})$, bounded by the lower and upper moving asymptotes L and U , which satisfy:

$$x^{\min} \leq L < U \leq x^{\max} \quad (4.60)$$

The moving asymptotes delimit the analytical domain of the approximation and act as curvature-control parameters. Their distance $U - L$ regulates the conditioning of the local convex approximation, while their position confines the admissible variation of the design variable to a regularized neighbourhood of the reference configuration.

The convex approximation of the nonlinear and nonconvex objective function is expressed in reciprocal form as:

$$\tilde{f}(x) = r + \frac{p}{U-x} + \frac{q}{x-L} \quad (4.61)$$

where $p > 0$, $q > 0$, and $r \in \mathbb{R}$ are the approximation parameters. The reciprocal structure in Equation (4.61) couples two monotonic terms with opposite trends: the term $q/(x-L)$ decreases on $(L, +\infty)$ and diverges to $+\infty$ as $x \rightarrow L^+$, whereas $p/(U-x)$ increases on $(-\infty, U)$ and diverges to $+\infty$ as $x \rightarrow U^-$. Their combination produces a smooth convex curve within (L, U) , whose curvature is governed by p and q , while r introduces a vertical translation that preserves convexity.

The first- and second-order derivatives of Equation (4.61) read:

$$\tilde{f}'(x) = \frac{p}{(U-x)^2} - \frac{q}{(x-L)^2} \quad \tilde{f}''(x) = \frac{2p}{(U-x)^3} + \frac{2q}{(x-L)^3} \quad (4.62)$$

Since $p, q > 0$, it follows that $\tilde{f}''(x) > 0$ for all $x \in (L, U)$, ensuring strict convexity of the approximation. The corresponding convex optimization problem formulated as:

$$\min_{x \in (L, U)} \tilde{f}(x) \quad (4.63)$$

admits a unique minimizer $x^* \in (L, U)$. The reciprocal terms act as analytical barriers that prevent minimizers from approaching the moving asymptotes interval, thereby stabilizing the curvature and conditioning of the local convex approximation.

The parameters p and q act as curvature weights controlling the steepness and symmetry of the approximation, while r guarantees value consistency with the nonlinear and nonconvex objective function $f(x)$.

From a theoretical standpoint, Equation (4.61) represents the minimal differentiable structure capable of ensuring strict convexity, local tunability, and boundedness within the moving asymptotes interval (L, U) .

4.3.2.1.2 Determination of Approximation Parameters

The local convex approximation of the objective function $\tilde{f}(x)$ introduced in Equation (4.61) is analytically determined so as to reproduce the local behaviour of the original nonlinear and nonconvex objective function $f(x)$ around a reference configuration x_0 strictly contained in the moving asymptotes interval, i.e., $L < x_0 < U$. The approximation parameters p , q , and r uniquely define the local convex approximation and must satisfy three analytical requirements: (i) first-order local consistency with the original nonlinear and nonconvex objective function, (ii) preservation of local monotonicity, and (iii) strict convexity over the moving asymptotes interval (L, U) .

Local consistency is enforced by requiring that both the value and the derivative of the approximation coincide with those of the original nonlinear and nonconvex objective function at x_0 :

$$\tilde{f}(x_0) = f(x_0) \quad \tilde{f}'(x_0) = f'(x_0) \quad (4.64)$$

Substituting Equation (4.61) into the first condition in Equation (4.64) yields:

$$r + \frac{p}{U - x_0} + \frac{q}{x_0 - L} = f(x_0) \quad (4.65)$$

from which the approximation parameter r is obtained as:

$$r = f(x_0) - \frac{p}{U - x_0} - \frac{q}{x_0 - L} \quad (4.66)$$

Differentiating Equation (4.61) and applying the second equation of Equation (4.64) provides the following linear relation between p and q :

$$\frac{p}{(U - x_0)^2} - \frac{q}{(x_0 - L)^2} = f'(x_0) \quad (4.67)$$

Equation (4.67) involves two unknowns and therefore requires an additional analytical condition to uniquely determine p and q . The additional requirement is introduced to preserve the local monotonicity of $f(x)$ at x_0 and to activate only the reciprocal term consistent with the sign of the local derivative $f'(x_0)$.

Specifically, for a locally increasing behaviour $f'(x_0) > 0$, the decreasing term $q/(x - L)$ is deactivate by setting $q = 0$, and Equation (4.67) gives:

$$p = (U - x_0)^2 f'(x_0) \quad (4.68)$$

whereas for a locally decreasing behaviour $f'(x_0) < 0$, the increasing term $p/(U - x)$ is deactivated by setting $p = 0$, which yields:

$$q = -(x_0 - L)^2 f'(x_0) \quad (4.69)$$




Figure 4.7: Monotonicity-preserving local convex MMA approximations obtained by selectively activating a single reciprocal term in the one-dimensional setting. Depending on the sign of the local derivative $f'(x_0)$, either the upper or the lower moving asymptote governs the curvature and numerical conditioning of the approximation, while first-order consistency with the original nonlinear and nonconvex objective function is preserved at the reference configuration x_0 . (a) Locally increasing objective function, with active upper asymptote ($p > 0$, $q = 0$). (b) Locally decreasing objective function, with active lower asymptote ($q > 0$, $p = 0$).

Therefore the approximation parameters p and q can be compactly written as:

$$p = \begin{cases} (U - x_0)^2 f'(x_0) & \text{if } f'(x_0) > 0 \\ 0 & \text{if } f'(x_0) \leq 0 \end{cases} \quad (4.70)$$

$$q = \begin{cases} 0 & \text{if } f'(x_0) \geq 0 \\ -(x_0 - L)^2 f'(x_0) & \text{if } f'(x_0) < 0 \end{cases} \quad (4.71)$$

The resulting approximation preserves the sign of the derivative and ensures that $p, q > 0$ whenever active, thereby maintaining strict convexity according to Equation (4.62). The

corresponding value of r follows from Equation (4.66), guaranteeing exact matching with $f(x_0)$. The parameter p or q governs the curvature and asymptotic behaviour of the approximation on the active side, while the inactive term is omitted to prevent distortion of local monotonicity.

4.3.2.1.3 Adaptive Update of the Moving Asymptotes

The moving asymptotes L and U define the analytical domain of the convex approximation $\tilde{f}(x)$ and govern its curvature behaviour. Their position determines the regularity and conditioning of the local approximation, and therefore represents a key element in the analytical structure of the MMA. Let $x_0 \in (L, U)$ denote the reference configuration at which the convex approximation is constructed. The moving asymptotes, L and U , must satisfy the condition in Equation (4.60) and remain separated from x_0 by strictly positive margins to avoid curvature singularities.

A basic configuration can be introduced by a symmetric offset around the reference configuration x_0 :

$$L = x_0 - \sigma \quad U = x_0 + \sigma \quad (4.72)$$

where $\sigma > 0$ is a scale factor consistent with the admissible design interval $[x^{\min}, x^{\max}]$. To guarantee a finite curvature and prevent ill-conditioning, the separation from the reference configuration must satisfy:

$$x_0 - L \geq \tau_{\min} > 0 \quad U - x_0 \geq \tau_{\min} > 0 \quad (4.73)$$

which provides a uniform positive curvature bound:

$$\tilde{f}''(x) \geq \frac{2 \min\{p, q\}}{\tau_{\min}^3} \quad \forall x \in (L, U) \quad (4.74)$$

ensuring strict convexity for $p, q > 0$.

The spacing of the moving asymptotes is adapted to the local variation of the original nonlinear and nonconvex objective function $f(x)$ through a non-negative scalar indicator $\delta(x_0)$, which quantifies the local variation scale. The adaptive update of the moving asymptotes rule is defined as:

$$L = x_0 - s \delta(x_0) \quad U = x_0 + s \delta(x_0) \quad (4.75)$$

where $s > 0$ is an adaptation coefficient that regulates the relationship between the moving asymptotes spacing and the local indicator $\delta(x_0)$. An increase in $\delta(x_0)$ extends the moving-asymptotes interval and reduces curvature, whereas a decrease in $\delta(x_0)$ narrows the moving asymptotes interval and increases curvature. The curvature of the approximation satisfies:

$$\frac{2 \min\{p, q\}}{(s \delta(x_0))^3} \leq \tilde{f}''(x_0) \leq \frac{2 \max\{p, q\}}{(s \delta(x_0))^3} \quad (4.76)$$

which confirms the role of $\delta(x_0)$ as an analytical curvature regulator.

Compatibility with the admissible design domain and non ill-conditioning of the moving asymptotes interval are maintained through the following conditions:

$$L = \max\{x^{\min}, x_0 - \max(\tau_{\min}, s \delta(x_0))\} \quad U = \min\{x^{\max}, x_0 + \max(\tau_{\min}, s \delta(x_0))\} \quad (4.77)$$

which guarantees the validity of Equations (4.60) and (4.73) for any admissible value of $\delta(x_0)$.

The coefficient s determines the rate of adaptation: small values generate narrower intervals and steeper curvature, while larger values produce broader intervals and milder curvature. In all cases, the positive lower bound in Equation (4.74) and the conditions in Equation (4.77) ensure analytical regularity, strict convexity, and bounded curvature within the moving asymptotes interval (L, U) .

From an analytical viewpoint, the relations in Equations (4.75) and (4.77) define a self-regulating mechanism that (i) maintains strict convexity of $\tilde{f}(x)$ within (L, U) ; (ii) prevents ill-conditioning by enforcing positive separations $x_0 - L$ and $U - x_0$; and (iii) aligns curvature control with the local variation of $f(x)$ through the indicator $\delta(x_0)$. The adaptive update of the moving asymptotes is consistent with the approximations parameters discussed in Section 4.3.2.1.2, since p and q influence curvature through the relations in Equation (4.76). The resulting convex approximation remains first-order consistent at x_0 (Equation (4.64)) and analytically conditioned across the interval (L, U) , providing a robust basis for the overall MMA.

4.3.2.1.4 Analytical Relation to Gradient-Based Methods

The MMA can be interpreted as a curvature-regularized extension of classical gradient-based formulations [263]. Conventional gradient-based approaches [233, 264, 265] rely on the local approximation of the nonlinear and nonconvex objective function, where the approximation is constructed through the first-order Taylor expansion around a reference configuration. In contrast, the MMA introduces a strictly convex reciprocal approximation that embeds curvature information explicitly through the analytical structure of the terms $q/(x - L)$ and $p/(U - x)$, as defined in Equations (4.70) and (4.71).

The analytical connection between MMA and classical gradient-based formulations is obtained by considering the limiting behaviour of the convex approximation in Equation (4.61). When the moving asymptotes are displaced toward infinity, i.e. $L \rightarrow -\infty$ and $U \rightarrow +\infty$, the reciprocal contributions vanish and the curvature of the approximation tends to zero. Therefore, the convex approximation reduces to the first-order Taylor expansion of the nonlinear and nonconvex objective function around the reference configuration x_0 :

$$\tilde{f}(x) = f(x_0) + f'(x_0)(x - x_0) \quad (4.78)$$

demonstrating that the MMA formulation analytically generalizes a first-order linear approximation of the classical gradient-based methods by introducing curvature modulation.

Equation (4.78) establishes the theoretical link between curvature control and first-order consistency. The convex approximation remains consistent with the local differential behaviour of $f(x)$ while incorporating an analytical mechanism that regularizes curvature through the finite positioning of the moving asymptotes. The resulting formulation unifies the local accuracy of method with the stability of convex representations, providing a rigorous analytical foundation for optimization in nonlinear and nonconvex settings [235, 236, 266, 267].

4.3.2.2 Extension to One-Dimensional Constrained Structural Optimization Problems

The extension of the theoretical foundations of the MMA to constrained structural optimization builds upon the one-dimensional unconstrained formulation introduced in Section 4.3.2.1 (see Equation (4.59)). The same one-dimensional design space and admissible bounds $[x^{\min}, x^{\max}]$ are retained, while a nonlinear and nonconvex inequality constraint function $g(x)$ is introduced. Accordingly, the constrained structural optimization problem reads:

$$\begin{aligned} \min_{x \in (x^{\min}, x^{\max})} \quad & f(x) \\ \text{subject to} \quad & g(x) \leq 0 \end{aligned} \quad (4.79)$$

where $f(x)$ denotes the original nonlinear and generally nonconvex objective function, and $g(x)$ is a nonlinear and nonconvex constraint function.

As in Section 4.3.2.1.1, the local convex approximation is constructed in a neighbourhood of a reference configuration $x_0 \in (x^{\min}, x^{\max})$ by introducing the moving asymptotes L and U , which satisfy the admissibility condition in Equation (4.60). Both $f(x)$ and $g(x)$ are represented by strictly convex approximations, $\tilde{f}(x)$ and $\tilde{g}(x)$, designed to ensure first-order consistency with the original problem in Equation (4.79), preserve local monotonicity, and maintain strict convexity within the moving-asymptote interval, as discussed in Section 4.3.1.3. Introducing the local convex approximations of the objective and constraint functions, the associated convex subproblem can be expressed as:

$$\begin{aligned} \min_{x \in (x^{\min}, x^{\max})} \quad & \tilde{f}(x) \\ \text{subject to} \quad & \tilde{g}(x) \leq 0. \end{aligned} \quad (4.80)$$

4.3.2.2.1 Convex Approximation of the Constraint Function

The construction of the convex approximation for the constraint function follows the same analytical procedure adopted for the objective function in Equation (4.61). Given the nonlinear and generally nonconvex constraint $g(x)$ introduced in Equation (4.79), a local convex approximation $\tilde{g}(x)$ is defined within the moving-asymptote interval $(L, U) \subset (x^{\min}, x^{\max})$ as:

$$\tilde{g}(x) = r_g + \frac{p_g}{U - x} + \frac{q_g}{x - L} \quad (4.81)$$

where $p_g > 0$, $q_g > 0$, and $r_g \in \mathbb{R}$ are the approximation parameters associated with the constraint function. The reciprocal structure in Equation (4.81) ensures smoothness and strict convexity over the analytical interval (L, U) , provided that the approximation parameters p_g and q_g satisfy the same positivity and separability conditions as in the objective function approximation.

The approximation parameters p_g , q_g , and r_g are determined to guarantee first-order consistency between the original constraint $g(x)$ and its convex approximation $\tilde{g}(x)$ at a reference configuration $x_0 \in (L, U)$. The consistency condition requires:

$$\tilde{g}(x_0) = g(x_0) \quad \tilde{g}'(x_0) = g'(x_0) \quad (4.82)$$

Substituting Equation (4.81) into the first relation of Equation (4.82) yields:

$$r_g = g(x_0) - \frac{p_g}{U - x_0} - \frac{q_g}{x_0 - L} \quad (4.83)$$

Differentiating Equation (4.81) and applying the second condition in Equation (4.82) leads to a linear relation between p_g and q_g :

$$\frac{p_g}{(U - x_0)^2} - \frac{q_g}{(x_0 - L)^2} = g'(x_0) \quad (4.84)$$

To uniquely determine p_g and q_g , the same monotonicity-preserving rule introduced in Section 4.3.2.1.2 is adopted. When the local derivative $g'(x_0) > 0$, the constraint function increases locally, and the decreasing reciprocal term $q_g/(x - L)$ is deactivated by setting $q_g = 0$; conversely, when $g'(x_0) < 0$, the increasing reciprocal term $p_g/(U - x)$ is deactivated by setting $p_g = 0$. The resulting expressions are:

$$p_g = \begin{cases} (U - x_0)^2 g'(x_0) & \text{if } g'(x_0) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (4.85)$$

$$q_g = \begin{cases} 0, & \text{if } g'(x_0) \geq 0 \\ -(x_0 - L)^2 g'(x_0) & \text{if } g'(x_0) < 0 \end{cases} \quad (4.86)$$

Substituting Equations (4.85) and (4.86) into Equation (4.83) provides the value of r_g , thus completing the definition of the convex approximation $\tilde{g}(x)$. The resulting function is continuously differentiable and strictly convex within the moving-asymptote interval (L, U) .

4.3.2.2.2 Lagrangian Formulation of the One-Dimensional Constrained Structural Optimization Problem

The extension of the MMA to nonlinear and nonconvex constrained structural optimization problems (see Equation (4.79)) requires a Lagrangian formulation that consistently incorporates inequality constraints while preserving the convex and separable structure of the convex subproblem defined in Equation (4.80). The Lagrangian formulation introduces an auxiliary functional that embeds both the objective and the constraint functions into a single analytical expression, establishing a direct relation between the primal and dual formulations of the problem. The primal formulation corresponds to the original minimization problem expressed in the variable x , whereas the dual formulation arises from the maximization with respect to the Lagrange multiplier λ , which enforces the constraint analytically. Such coupling provides a unified analytical representation that ensures convexity, differentiability, and feasibility within the MMA formulation.

The Lagrangian functional associated with the convex subproblem is defined as:

$$\mathcal{L}(x, \lambda) = \tilde{f}(x) + \lambda \tilde{g}(x) \quad \lambda \geq 0 \quad (4.87)$$

where $\tilde{f}(x)$ and $\tilde{g}(x)$ denote the convex and continuously differentiable approximations of the nonlinear and nonconvex functions $f(x)$ and $g(x)$ introduced in Equations (4.61) and (4.81), and λ represents the Lagrange multiplier associated with the inequality constraint. The variable x defines the primal space, containing the feasible configurations of the design variable within the analytical domain (L, U) , while λ defines the dual space, associated with the analytical enforcement of the constraint. Minimization of $\mathcal{L}(x, \lambda)$ with respect to x identifies the optimal primal configuration, whereas maximization with respect to $\lambda \geq 0$ enforces constraint satisfaction through the dual variable.

Substituting the reciprocal representations of $\tilde{f}(x)$ and $\tilde{g}(x)$ into Equation (4.87) yields

$$\mathcal{L}(x, \lambda) = (r + \lambda r_g) + \frac{p + \lambda p_g}{U - x} + \frac{q + \lambda q_g}{x - L} : \quad (4.88)$$

which is continuously differentiable with respect to both x and λ in the analytical domain $(L, U) \times [0, \infty)$. The constrained minimization problem is thereby transformed into a saddle-point problem, where the primal variable x minimizes $\mathcal{L}(x, \lambda)$ over (L, U) , and the dual variable λ maximizes it over $[0, \infty)$. Such structure ensures a consistent analytical correspondence between the primal and dual representations and preserves the convexity and separability properties of the local subproblem.

Under the assumptions of strict convexity and differentiability of $\tilde{f}(x)$ and $\tilde{g}(x)$, the necessary and sufficient first-order optimality conditions are given by the Karush–Kuhn–Tucker (KKT) conditions:

$$\frac{\partial \mathcal{L}}{\partial x}(x^*, \lambda^*) = 0 \quad (4.89)$$

$$\tilde{g}(x^*) \leq 0 \quad (4.90)$$

$$\lambda^* \geq 0 \quad (4.91)$$

$$\lambda^* \tilde{g}(x^*) = 0 \quad (4.92)$$

where Equation (4.89) enforces stationarity of the Lagrangian with respect to the primal variable, Equation (4.90) defines primal feasibility, Equation (4.91) imposes dual feasibility, and Equation (4.92) expresses complementary slackness. In the case of inequality constraints, the set of conditions Equations (4.90) to (4.92) replaces the standard derivative with respect to λ used in equality-constrained formulations, ensuring a proper analytical characterization of the constrained equilibrium.

Differentiation of Equation (4.88) with respect to x yields:

$$\frac{\partial \mathcal{L}}{\partial x} = \frac{p}{(U - x)^2} - \frac{q}{(x - L)^2} + \lambda \left(\frac{p_g}{(U - x)^2} - \frac{q_g}{(x - L)^2} \right) \quad (4.93)$$

and the condition $\partial \mathcal{L}(x^*, \lambda^*) / \partial x = 0$ represents the equilibrium between the convex approximations of the objective and constraint functions at the optimal configuration (x^*, λ^*) . If the constraint is inactive, the feasibility condition in Equation (4.90) implies $\tilde{g}(x^*) < 0$, and

complementary slackness in Equation (4.92) enforces $\lambda^* = 0$, leading to the unconstrained minimizer of $\tilde{f}(x)$. If the constraint is active, Equations (4.89) and (4.92) require $\tilde{g}(x^*) = 0$ and $\lambda^* > 0$, yielding the coupled system:

$$\frac{\partial \mathcal{L}}{\partial x} = 0 \quad \tilde{g}(x) = 0 \quad x \in (L, U) \quad \lambda > 0 \quad (4.94)$$

whose strict convexity guarantees existence and uniqueness of the solution within the analytical domain.

4.3.2.3 Extension to Multidimensional and Large-Scale Problems

The one-dimensional analytical formulation developed in Sections 4.3.2.1 and 4.3.2.2 establishes the theoretical foundations of the MMA. However, structural optimization problems typically involve multidimensional design spaces, where each design variable contributes independently to the global response, while multiple nonlinear and nonconvex inequality constraints collectively define the feasible domain. Extending the MMA to multidimensional problems requires the analytical generalization of the convex approximation, the preservation of convexity and separability across all design variables, and the implementation of explicit curvature-control mechanisms to ensure stability and regularity in large-scale settings.

4.3.2.3.1 Analytical Generalization and Separable Structure

The analytical generalization of the MMA to multidimensional constrained structural optimization problems extends the one-dimensional formulation introduced in Sections 4.3.2.1 and 4.3.2.2. Let the design vector be denoted by $\mathbf{x} = [x_1, x_2, \dots, x_A]^T \in \mathbb{R}_{\text{adm}}^A$, where A denotes the number of design variables and the admissible bounds are defined in Equation (4.50). After enforcing the discrete equilibrium equations $\mathbf{h}(\mathbf{x}) = \mathbf{0}$ at the finite element analysis level, the reduced nonlinear and nonconvex constrained structural optimization problem reads:

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}_{\text{adm}}^A} \quad & f(\mathbf{x}) \\ \text{subject to} \quad & \mathbf{g}(\mathbf{x}) \leq \mathbf{0} \end{aligned} \quad (4.95)$$

where $f(\mathbf{x}) : \mathbb{R}^A \rightarrow \mathbb{R}$ denotes the objective function and $\mathbf{g}(\mathbf{x}) : \mathbb{R}^A \rightarrow \mathbb{R}^C$ is the vector of nonlinear and nonconvex inequality constraints.

The local convex approximation is constructed around a reference configuration $\mathbf{x}_0 \in \mathbb{R}_{\text{adm}}^A$ by introducing, for each design variable x_a , a couple of moving asymptotes L_a and U_a satisfying:

$$x_a^{\min} \leq L_a < x_{0a} < U_a \leq x_a^{\max} \quad a = 1, \dots, A \quad (4.96)$$

Each couple (L_a, U_a) defines the analytical interval for the local convex approximation along direction x_a and provides explicit control of the local curvature.

The nonlinear and nonconvex objective function $f(\mathbf{x})$ is approximated by the separable convex function:

$$\tilde{f}(\mathbf{x}) = r + \sum_{a=1}^A \left(\frac{p_a}{U_a - x_a} + \frac{q_a}{x_a - L_a} \right) \quad (4.97)$$

where $p_a > 0$ and $q_a > 0$ (when active) regulate the curvature of the local convex approximation along coordinate x_a , and $r \in \mathbb{R}$ enforces interpolation of the function value at \mathbf{x}_0 . The approximation in Equation (4.97) ensures smoothness and strict convexity within the analytical domain $(L_1, U_1) \times \cdots \times (L_A, U_A)$.

The approximation parameters p_a , q_a , and r are determined to guarantee first-order consistency between the nonlinear and nonconvex function $f(\mathbf{x})$ and its convex approximation $\tilde{f}(\mathbf{x})$ at the reference configuration \mathbf{x}_0 :

$$\tilde{f}(\mathbf{x}_0) = f(\mathbf{x}_0) \quad \nabla_{\mathbf{x}} \tilde{f}(\mathbf{x}_0) = \nabla_{\mathbf{x}} f(\mathbf{x}_0) \quad (4.98)$$

From Equation (4.97) and Equation (4.98) one obtains:

$$r = f(\mathbf{x}_0) - \sum_{a=1}^A \left(\frac{p_a}{U_a - x_{0a}} + \frac{q_a}{x_{0a} - L_a} \right) \quad (4.99)$$

and:

$$\frac{p_a}{(U_a - x_{0a})^2} - \frac{q_a}{(x_{0a} - L_a)^2} = \frac{\partial f}{\partial x_a}(\mathbf{x}_0) = f_{x_a}(\mathbf{x}_0) \quad a = 1, \dots, A \quad (4.100)$$

To preserve local monotonicity, the same rule adopted in Section 4.3.2.1.2 is applied. When the partial derivative $f_{x_a}(\mathbf{x}_0) > 0$, the decreasing reciprocal term $q_a/(x_a - L_a)$ is deactivated by setting $q_a = 0$; conversely, when $f_{x_a}(\mathbf{x}_0) < 0$, the increasing reciprocal term $p_a/(U_a - x_a)$ is deactivated by setting $p_a = 0$. The resulting expressions read:

$$p_a = \begin{cases} (U_a - x_{0a})^2 f_{x_a}(\mathbf{x}_0) & \text{if } f_{x_a}(\mathbf{x}_0) > 0 \\ 0, & \text{if } f_{x_a}(\mathbf{x}_0) \leq 0 \end{cases} \quad (4.101)$$

$$q_a = \begin{cases} 0, & \text{if } f_{x_a}(\mathbf{x}_0) \geq 0 \\ -(x_{0a} - L_a)^2 f_{x_a}(\mathbf{x}_0) & \text{if } f_{x_a}(\mathbf{x}_0) < 0 \end{cases} \quad (4.102)$$

If $f_{x_a}(\mathbf{x}_0) = 0$, curvature regularization is introduced by prescribing a strictly positive second derivative $\tilde{f}_{x_a x_a}(\mathbf{x}_0) = \kappa_a > 0$, distributed between p_a and q_a by a weight parameter $\alpha_a \in (0, 1)$:

$$\frac{2p_a}{(U_a - x_{0a})^3} = \alpha_a \kappa_a, \quad \frac{2q_a}{(x_{0a} - L_a)^3} = (1 - \alpha_a) \kappa_a. \quad (4.103)$$

Each component of the nonlinear and nonconvex constraint vector $\mathbf{g}(\mathbf{x})$ is represented by a local convex approximation of the form:

$$\tilde{\mathbf{g}}(\mathbf{x}) = \mathbf{r}_{\mathbf{g}} + \sum_{a=1}^A \left(\frac{\mathbf{p}_{\mathbf{g},a}}{U_a - x_a} + \frac{\mathbf{q}_{\mathbf{g},a}}{x_a - L_a} \right) \in \mathbb{R}^C \quad (4.104)$$

where $\mathbf{p}_{\mathbf{g},a}, \mathbf{q}_{\mathbf{g},a} \in \mathbb{R}_{\geq 0}^C$ collect the approximation parameters of the constraint functions along direction x_a , and $\mathbf{r}_{\mathbf{g}} \in \mathbb{R}^C$ enforces interpolation at \mathbf{x}_0 . The first-order consistency conditions are:

$$\tilde{\mathbf{g}}(\mathbf{x}_0) = \mathbf{g}(\mathbf{x}_0) \quad \nabla_{\mathbf{x}} \tilde{\mathbf{g}}(\mathbf{x}_0) = \nabla_{\mathbf{x}} \mathbf{g}(\mathbf{x}_0) \quad (4.105)$$

which imply:

$$\mathbf{r}_g = \mathbf{g}(\mathbf{x}_0) - \sum_{a=1}^A \left(\frac{\mathbf{p}_{g,a}}{U_a - x_{0a}} + \frac{\mathbf{q}_{g,a}}{x_{0a} - L_a} \right) \quad (4.106)$$

and, for each constraint g_c :

$$\frac{p_{g_c,a}}{(U_a - x_{0a})^2} - \frac{q_{g_c,a}}{(x_{0a} - L_a)^2} = \frac{\partial g_c}{\partial x_a}(\mathbf{x}_0) = g_{c,x_a}(\mathbf{x}_0) \quad (4.107)$$

The same monotonicity-preserving rule in Equations (4.101) and (4.102) is applied to $p_{g_c,a}$ and $q_{g_c,a}$ to ensure strict convexity of each component of the constraint approximation.

Since the convex approximations in Equations (4.97) and (4.104) are separable with respect to the design variables, the Hessian of $\tilde{f}(\mathbf{x})$ is diagonal:

$$\nabla_{\mathbf{x}}^2 \tilde{f}(\mathbf{x}) = \text{diag} \left(\frac{2p_1}{(U_1 - x_1)^3} + \frac{2q_1}{(x_1 - L_1)^3}, \dots, \frac{2p_A}{(U_A - x_A)^3} + \frac{2q_A}{(x_A - L_A)^3} \right) \quad (4.108)$$

and is strictly positive definite in the analytical domain whenever $p_a, q_a > 0$. The same argument applies to each scalar constraint $\tilde{g}_c(\mathbf{x})$, which is strictly convex in \mathbf{x} .

Accordingly, the multidimensional convex subproblem reads:

$$\begin{aligned} & \min_{\mathbf{x} \in \mathbb{R}_{\text{adm}}^A} \tilde{f}(\mathbf{x}) \\ & \text{subject to } \tilde{\mathbf{g}}(\mathbf{x}) \leq \mathbf{0} \end{aligned} \quad (4.109)$$

which defines a strictly convex and separable problem that admits a unique minimizer within the analytical domain. The moving asymptotes control the curvature and the analytical conditioning along each coordinate direction and are maintained strictly separated from \mathbf{x}_0 :

$$x_{0a} - L_a \geq \tau_{\min,a} > 0 \quad U_a - x_{0a} \geq \tau_{\min,a} > 0 \quad a = 1, \dots, A \quad (4.110)$$

Under these conditions, the multidimensional MMA yields a strictly convex, separable, and first-order consistent local convex approximation of the nonlinear and nonconvex structural optimization problem in Equation (4.95), with explicit curvature control and analytical regularity provided by the moving asymptotes. The analytical structure presented in this section provides the foundation for the algorithmic implementation of the MMA, discussed in the subsequent section.

4.3.3 Algorithmic Structure and Computational Implementation

The analytical development in Section 4.3.2 establishes the theoretical foundations of the MMA. The present section formalizes the algorithmic structure and its computational implementation with emphasis on the iterative construction of convex subproblems, the Lagrangian formulation, and the numerical update of the design variables, the moving asymptotes, and approximation parameters. The objective is to connect the analytical expressions derived in Section 4.3.2 with the computational operations required for a robust and reproducible implementation.

4.3.3.1 Lagrangian Formulation and Optimality Conditions

The solution of the convex subproblem is obtained through its Lagrangian formulation, which provides the analytical foundation for deriving the optimality conditions governing the primal and dual problems as discussed in Section 4.3.2.2. Let $\boldsymbol{\lambda}^{(k)} = [\lambda_1^{(k)}, \dots, \lambda_J^{(k)}]^\top \geq \mathbf{0}$ denote the vector of Lagrange multipliers associated with the inequality constraints $\tilde{t}_j^{(k)}(\mathbf{x}) \leq 0$. The primal variables are represented by the vector of design variables \mathbf{x} , bounded within the move limits $\boldsymbol{\alpha}^{(k)} \leq \mathbf{x} \leq \boldsymbol{\beta}^{(k)}$, while the dual variables are represented by the Lagrange multipliers $\boldsymbol{\lambda}$. The Lagrangian function of the convex subproblem at iteration k is defined as:

$$\mathcal{L}^{(k)}(\mathbf{x}, \boldsymbol{\lambda}^{(k)}) = \tilde{t}_0^{(k)}(\mathbf{x}) + \sum_{j=1}^J \lambda_j^{(k)} \tilde{t}_j^{(k)}(\mathbf{x}) \quad (4.111)$$

which combines the primal and dual variables into a single scalar function representing the augmented form of the convex subproblem.

The primal problem seeks the optimal design configuration that minimizes the local convex approximation of the objective function while satisfying the inequality and move limits constraints:

$$\begin{aligned} \min_{\boldsymbol{\alpha}^{(k)} \leq \mathbf{x} \leq \boldsymbol{\beta}^{(k)}} \quad & \tilde{t}_0^{(k)}(\mathbf{x}) \\ \text{subject to} \quad & \tilde{t}_j^{(k)}(\mathbf{x}) \leq 0, \quad j = 1, \dots, J \end{aligned} \quad (4.112)$$

The associated dual function is obtained by minimizing the Lagrangian function with respect to the primal variables \mathbf{x} :

$$\varphi^{(k)}(\boldsymbol{\lambda}) = \min_{\boldsymbol{\alpha}^{(k)} \leq \mathbf{x} \leq \boldsymbol{\beta}^{(k)}} \mathcal{L}^{(k)}(\mathbf{x}, \boldsymbol{\lambda}) \quad (4.113)$$

The corresponding dual problem consists of maximizing the dual function with respect to the dual variables $\boldsymbol{\lambda}$:

$$\max_{\boldsymbol{\lambda} \geq \mathbf{0}} \varphi^{(k)}(\boldsymbol{\lambda}) \quad (4.114)$$

The convexity of the subproblem guarantees that the optimal solutions of the primal and dual problems coincide, and that the primal and dual variables, $\mathbf{x}^{(k+1)}$ and $\boldsymbol{\lambda}^{(k)}$, satisfies the necessary and sufficient optimality conditions.

The first-order optimality conditions for the convex subproblem correspond to the Karush–Kuhn–Tucker (KKT) system:

$$\nabla_{\mathbf{x}} \mathcal{L}^{(k)}(\mathbf{x}^{(k+1)}, \boldsymbol{\lambda}^{(k)}) = \mathbf{0} \quad (\text{stationarity}) \quad (4.115)$$

$$\tilde{t}_j^{(k)}(\mathbf{x}^{(k+1)}) \leq 0 \quad (\text{primal feasibility}) \quad (4.116)$$

$$\boldsymbol{\alpha}^{(k)} \leq \mathbf{x}^{(k+1)} \leq \boldsymbol{\beta}^{(k)} \quad (\text{box feasibility}) \quad (4.117)$$

$$\lambda_j^{(k)} \geq 0 \quad (\text{dual feasibility}) \quad (4.118)$$

$$\lambda_j^{(k)} \tilde{t}_j^{(k)}(\mathbf{x}^{(k+1)}) = 0 \quad (\text{complementarity}) \quad (4.119)$$

The first condition enforces the vanishing of the gradient of the Lagrangian function with respect to the primal variables, while the remaining relations ensure feasibility and complementarity between the primal and dual variables.

Due to the separable structure of the convex approximation, the stationarity condition in Equation (4.115) can be written in closed form for each design variable. Differentiating the Lagrangian function with respect to x_a yields:

$$\frac{\partial \mathcal{L}^{(k)}}{\partial x_a} = \frac{P_a^{(k)}}{(U_a^{(k)} - x_a)^2} - \frac{Q_a^{(k)}}{(x_a - L_a^{(k)})^2} \quad (4.120)$$

where the combined approximation parameters are defined as:

$$P_a^{(k)} = p_{a0}^{(k)} + \sum_{j=1}^J \lambda_j^{(k)} p_{aj}^{(k)} \quad Q_a^{(k)} = q_{a0}^{(k)} + \sum_{j=1}^J \lambda_j^{(k)} q_{aj}^{(k)} \quad (4.121)$$

Imposing the stationarity condition $\partial \mathcal{L}^{(k)} / \partial x_a = 0$ for $\alpha_a^{(k)} < x_a < \beta_a^{(k)}$ leads to the analytical expression of the stationary solution:

$$x_a^* = \frac{\sqrt{P_a^{(k)} L_a^{(k)}} + \sqrt{Q_a^{(k)} U_a^{(k)}}}{\sqrt{P_a^{(k)}} + \sqrt{Q_a^{(k)}}} \quad (4.122)$$

which is well defined for $P_a^{(k)} > 0$ and $Q_a^{(k)} > 0$, conditions ensured by the convexity of the local approximation and by the nonnegativity of the Lagrange multipliers. The stationary solution is then projected into the admissible interval to preserve feasibility with respect to the move limits:

$$x_a^{(k+1)} = \min\{\beta_a^{(k)}, \max\{\alpha_a^{(k)}, x_a^*\}\} \quad (4.123)$$

which enforces $\alpha^{(k)} \leq \mathbf{x}^{(k+1)} \leq \beta^{(k)}$ and ensures the satisfaction of the KKT conditions.

For a fixed vector of dual variables λ , the minimization of the Lagrangian function in Equation (4.113) is separable with respect to the primal variables and admits the analytical solution given by Equations (4.122) and (4.123). Substituting the stationary solution into the Lagrangian function yields the corresponding dual function:

$$\varphi^{(k)}(\lambda) = \mathcal{L}^{(k)}(\mathbf{x}^{(k+1)}(\lambda), \lambda) \quad (4.124)$$

whose gradient with respect to the dual variables is:

$$\frac{\partial \varphi^{(k)}}{\partial \lambda_j} = \tilde{f}_j^{(k)}(\mathbf{x}^{(k+1)}(\lambda)) \quad j = 1, \dots, J \quad (4.125)$$

The dual variables $\lambda^{(k)}$ are obtained by solving the dual problem (4.114), typically through projected Newton iterations or gradient-based updates constrained to $\lambda \geq \mathbf{0}$, until both dual feasibility and complementarity are satisfied.

At convergence, $\mathbf{x}^{(k+1)}$, $\lambda^{(k)}$ satisfies the complete set of KKT conditions Equations (4.115) and (4.119) and defines the unique primal–dual solution of the convex subproblem at iteration k . The Lagrangian formulation therefore provides a rigorous and computationally consistent structure for coupling the update of the design variables with the enforcement of constraint feasibility through the Lagrange multipliers. The separable nature of the local convex approximation ensures analytical tractability, strict convexity, and numerical robustness, thereby forming the computational basis for the iterative implementation of the MMA.

4.3.3.2 Adaptive Update of the Moving Asymptotes

The analytical formulation derived in Section 4.3.3.1 requires adaptive mechanisms to update the analytical domain of each design variable and maintain stability across iterations. The update strategy regulates the curvature of the local convex approximations and controls the numerical conditioning of the local subproblem by adjusting the analytical domain of each design variable according to its recent evolution.

At iteration k , the moving asymptotes are updated as:

$$L_a^{(k)} = x_a^{(k)} - s_a^{(k)}(x_a^{(k-1)} - L_a^{(k-1)}), \quad U_a^{(k)} = x_a^{(k)} + s_a^{(k)}(U_a^{(k-1)} - x_a^{(k-1)}), \quad (4.126)$$

where $L_a^{(k)}$ and $U_a^{(k)}$ denote the lower and upper moving asymptotes of the design variable x_a , and $s_a^{(k)}$ is a scaling coefficient governing the contraction or expansion of the moving-asymptote interval.

The coefficient $s_a^{(k)}$ depends on the sign of consecutive design-variable variations and is defined as:

$$s_a^{(k)} = \begin{cases} 0.7 & \text{if } (x_a^{(k)} - x_a^{(k-1)})(x_a^{(k-1)} - x_a^{(k-2)}) < 0, \\ 1.2 & \text{if } (x_a^{(k)} - x_a^{(k-1)})(x_a^{(k-1)} - x_a^{(k-2)}) > 0, \\ 1.0 & \text{otherwise.} \end{cases} \quad (4.127)$$

A contraction of the moving-asymptote interval ($s_a^{(k)} < 1$) occurs when consecutive design-variable updates exhibit oscillatory behavior, thereby enhancing stability and damping such oscillations. Conversely, an expansion ($s_a^{(k)} > 1$) accelerates convergence in regions characterized by monotonic evolution of the design variable. These two characteristic situations are schematically illustrated in Figure 4.8, which highlights the role of the adaptive scaling coefficient in regulating the local curvature and stability of the approximation.

To prevent curvature singularities and preserve the analytical regularity of the convex approximation, a minimum separation is enforced between each design variable and its moving asymptotes:

$$x_a^{(k)} - L_a^{(k)} \geq \tau_{\min,a} > 0, \quad U_a^{(k)} - x_a^{(k)} \geq \tau_{\min,a} > 0, \quad (4.128)$$

where $\tau_{\min,a}$ is a positive tolerance ensuring that the moving asymptotes remain at a finite distance from the current design configuration.




Figure 4.8: Adaptive update of the moving asymptotes in MMA. The scaling coefficient $s_a^{(k)}$ dynamically modifies the asymptote interval based on the local behavior of the design variable. (a) Oscillatory evolution across successive iterations ($x^{k-2} \rightarrow x^{k-1} \rightarrow x^k$), inducing a contraction of the asymptote interval to enhance numerical stability. (b) Monotonic evolution of the design variable, leading to an expansion of the asymptote interval and promoting faster convergence of the iterative process.

In parallel, the move limits $\alpha^{(k)}$ and $\beta^{(k)}$ define the admissible variation range for the design variables within the domain of admissibility of the convex approximation. The move limits are updated at each iteration as:

$$\alpha_a^{(k)} = \max(x_a^{\min}, L_a^{(k)} + \omega(x_a^{(k)} - L_a^{(k)}), x_a^{(k)} - \Delta x(x_a^{\max} - x_a^{\min})), \quad (4.129)$$

$$\beta_a^{(k)} = \min(x_a^{\max}, U_a^{(k)} - \omega(U_a^{(k)} - x_a^{(k)}), x_a^{(k)} + \Delta x(x_a^{\max} - x_a^{\min})), \quad (4.130)$$

where ω and Δx are dimensionless control parameters regulating the aggressiveness and amplitude of the move limits. Typical values used in structural-optimization applications are $\omega = 0.1$ and $\Delta x = 0.2$ [4, 33, 237–239].

The move limits ensure that the admissible interval remains strictly contained within the moving-asymptote interval, $L_a^{(k)} < \alpha_a^{(k)} \leq \beta_a^{(k)} < U_a^{(k)}$, thereby preserving the analytical consistency of the convex model and preventing infeasible updates of the design variables.

The adaptive update of the moving asymptotes and move limits therefore governs the balance between numerical stability and convergence rate, ensuring that the iterative process remains well-conditioned, stable, and robust across successive iterations of the MMA.

4.3.3.3 Iterative Procedure and Convergence

The adaptive mechanisms introduced in Section 4.3.3.2 are embedded within the overall iterative scheme of the MMA. At each iteration k , the procedure alternates between the construction of the separable convex subproblem, the solution of its Lagrangian formulation, and the adaptive update of the moving asymptotes and move limits. The algorithm proceeds until the updates of the design variables and the variations of the objective function satisfy prescribed stopping criteria while maintaining constraint feasibility within tolerances. The complete procedure is summarized in Algorithm 4, which consists of an outer loop on the iteration index k and an inner dual solve devoted to minimizing the Lagrangian with respect to the primal variables and maximizing it with respect to the dual variables.

From a computational standpoint, the inner dual solve exploits the separability of the Lagrangian function with respect to the design variables. For any fixed vector of Lagrange multipliers, the stationarity condition in Equation (4.115) reduces to Equation (4.120) with the closed-form stationary solutions in Equation (4.122), followed by the projection in Equation (4.123). The Lagrange multipliers are then updated via a projected ascent step on the dual function $\varphi^{(k)}$, using either a safeguarded gradient iteration or a projected Newton step. This nested scheme yields $\mathbf{x}^{(k+1)}$ and $\lambda^{(k)}$ that satisfy the KKT conditions Equations (4.115) and (4.119) within the prescribed tolerance ε_c .

The outer loop advances the design variables by reconstructing, at each iteration, the local convex approximations around $\mathbf{x}^{(k)}$, solving the corresponding convex subproblem, and adapting the moving asymptotes and move limits according to Equations (4.126), (4.127), (4.129) and (4.130). The stopping criteria simultaneously monitor the normalized change in the design variables, the normalized variation of the objective function, and the maximum constraint violation. When all measures fall below ε_x , ε_f , and ε_c , the sequence $\{\mathbf{x}^{(k)}\}$ is deemed to have reached a stationary configuration of the original constrained structural optimization problem within the accuracy specified by the tolerances. The safeguards on the asymptote scaling and on the amplitude of the move limits balance robustness and efficiency by damping oscillations and preventing ill-conditioning, while preserving the strict convexity and separability that ensure the computational tractability of the method.

4.3.4 Numerical Validation and Analysis of Algorithmic Behaviour

The present section provides a numerical validation of the MMA and an analysis of its algorithmic behaviour on representative structural optimization problems.

In accordance with the structure of the original manuscript [4], the remainder of the present section is organized as follows: Section 4.3.4.1 presents the results for the cantilever beam problem, Section 4.3.4.2 reports the solution of the 8-bar truss, and Section 4.3.4.3 investigates the 2-bar truss configuration. In each case, the results are used to illustrate the behavior of the algorithm, analyze its convergence properties, and assess the robustness and effectiveness of the MMA in solving non-linear and constrained structural optimization problems. All test cases are implemented in the MATHEMATICA[®] (Wolfram Research, Inc.) environment, consistently with

Algorithm 4: Method of Moving Asymptotes (MMA)

Input: initial design $\mathbf{x}^{(0)}$, bounds $\mathbf{x}^{\min}, \mathbf{x}^{\max}$, tolerances $\varepsilon_x, \varepsilon_f, \varepsilon_c$, parameters $\omega, \Delta x$
Output: converged design $\mathbf{x}^{(k)}$

- 1 **Initialization:** Choose $\mathbf{L}^{(0)}, \mathbf{U}^{(0)}$ around $\mathbf{x}^{(0)}$; set $\alpha^{(0)}, \beta^{(0)}$ (Eqs. (4.129)–(4.130)); $k \leftarrow 0$.
- 2 **while not converged do**
- 3 **Model construction:** Evaluate $f(\mathbf{x}^{(k)}), g_j(\mathbf{x}^{(k)})$ and gradients; build separable convex approximations $\tilde{t}_j^{(k)}$ (Eq. (4.128)).
- 4 **Subproblem:** Form convex subproblem (4.52) on $[\alpha^{(k)}, \beta^{(k)}]$.
- 5 **Dual solution:** Initialize $\lambda^{(k)} \geq 0$. **repeat**
- 6 | dual feasibility & complementarity $\leq \varepsilon_c$
- 7 **until** Compute $P_a^{(k)}, Q_a^{(k)}$ (Eq. (4.121)); compute x_a^* (Eq. (4.122)); project x_a^* (Eq. (4.123)); update $\lambda^{(k)}$;;
- 8 **Update:** $\mathbf{x}^{(k+1)} = \mathbf{x}^*$.
- 9 Update $\mathbf{L}^{(k+1)}, \mathbf{U}^{(k+1)}$ (Eqs. (4.126)–(4.127)).
- 10 Update $\alpha^{(k+1)}, \beta^{(k+1)}$ (Eqs. (4.129)–(4.130)), ensuring $L_a^{(k+1)} < \alpha_a^{(k+1)} \leq \beta_a^{(k+1)} < U_a^{(k+1)}$.
- 11 **Convergence:** Stop if
$$\frac{\|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\|_\infty}{1 + \|\mathbf{x}^{(k)}\|_\infty} \leq \varepsilon_x, \quad \frac{|f(\mathbf{x}^{(k+1)}) - f(\mathbf{x}^{(k)})|}{1 + |f(\mathbf{x}^{(k)})|} \leq \varepsilon_f, \quad \max_j (\tilde{t}_j^{(k)}(\mathbf{x}^{(k+1)}))_+ \leq \varepsilon_c.$$
- 12 **Safeguard:** If needed, reduce $s_a^{(k+1)}$ and Δx .
- 13 $k \leftarrow k + 1$;

the computational framework described in the introduction, ensuring full analytical control and numerical reproducibility.

4.3.4.1 Cantilever Beam with Displacement Constraint

The first benchmark problem concerns the structural optimization of a cantilever beam composed of five beam elements connected at six nodes, as illustrated in Figure 4.9. The beam is clamped at node 1 and subjected to a vertical point load at node 6. Each element is modeled as a prismatic beam with a square cross-section of fixed wall thickness and variable height x_a , which represents the design variable associated with the a -th element ($a = 1, \dots, 5$). The structural response is evaluated under linear elastic assumptions according to standard beam theory.

The optimization objective is to minimize the total structural weight, which is proportional to the sum of the cross-section heights x_a . A single nonlinear constraint is imposed to limit the vertical displacement at the free end (node 6) to a prescribed maximum value. The problem is




Figure 4.9: Cantilever beam configuration used in Test Problem 1. The structure consists of five beam elements connected at six nodes. Node 1 is clamped, while a vertical point load is applied at node 6. Each element has a square cross-section of variable height x_a and fixed wall thickness.

formulated as:

$$\begin{aligned} & \min_{x_a > 0} C_1 (x_1 + x_2 + x_3 + x_4 + x_5) \\ & \text{subject to } \frac{61}{x_1^3} + \frac{37}{x_2^3} + \frac{19}{x_3^3} + \frac{7}{x_4^3} + \frac{1}{x_5^3} \leq C_2 \end{aligned} \quad (4.131)$$

where $C_1 > 0$ is a constant depending on the material density and geometric properties of the cross-section, and $C_2 > 0$ denotes the maximum admissible displacement at the free end, obtained analytically from beam theory.

The problem is solved using different values of the initial moving-asymptote spacing parameter t , in accordance with the settings reported in [4]. The numerical results are summarized in Table 4.8, where each entry reports the total weight and the corresponding constraint violation at successive iterations for a given value of t .

The results confirm the correctness of the implementation, as the optimal values computed match those presented in the original reference. Moreover, the numerical trends highlight the influence of the moving asymptote spacing on the convergence characteristics of the algorithm. Small values of t produce narrow approximation intervals, which increase the curvature of the local convex approximation and yield conservative design updates. In such cases, the convergence is more stable but typically requires a larger number of iterations. Conversely, moderate values of t (e.g., $t = \frac{1}{3}$) allow wider analytical intervals, resulting in faster progress toward feasibility and optimality. However, excessively large values of t reduce curvature control and may induce oscillatory or erratic updates.

The obtained results are in full agreement with the observations of [4], emphasizing the critical role of the moving asymptote spacing in balancing stability and convergence rate. In particular, they confirm that the moving asymptotes effectively regulate the domain at each iteration and demonstrate the capability of the MMA to handle highly nonlinear and constrained structural optimization problems with accuracy and robustness.

4.3.4.2 Eight-Bar Truss under Stress Constraints

The second benchmark problem considers a three-dimensional truss structure composed of eight bars symmetrically connected to a central node, as illustrated in Figure 4.10. The external load is

Table 4.8: MMA results for the cantilever beam problem. Each entry reports the pair (weight, infeasibility) at each iteration for different asymptote-spacing parameters t .

Iter.	$t = 0$	$t = \frac{1}{16}$	$t = \frac{1}{8}$	$t = \frac{1}{4}$	$t = \frac{1}{3}$	$t = \frac{1}{2}$	$t = \frac{2}{3}$	$t = \frac{3}{4}$
0	1.560/0.000	1.560/0.000	1.560/0.000	1.560/0.000	1.560/0.000	1.560/0.000	1.560/0.000	1.560/0.000
1	1.265/0.400	1.274/0.350	1.285/0.230	1.309/0.100	1.327/0.050	1.387/0.000	1.448/0.000	1.477/0.000
2	1.251/0.430	1.270/0.270	1.307/0.110	1.335/0.010	1.338/0.004	1.346/0.000	1.386/0.000	1.418/0.000
3	1.259/0.430	1.304/0.140	1.331/0.030	1.340/0.0005	1.340/0.0001	1.341/0.000	1.358/0.000	1.383/0.000
4	1.250/0.440	1.319/0.080	1.337/0.008	—	—	1.347/0.000	1.363/0.000	—
5	1.258/0.430	1.329/0.040	1.339/0.002	—	—	1.343/0.000	1.352/0.000	—
6	1.249/0.440	1.333/0.020	1.340/0.001	—	—	1.341/0.000	1.346/0.000	—
7	1.258/0.430	1.336/0.010	—	—	—	—	1.343/0.000	—
8	—	—	—	—	—	—	1.342/0.000	—
9	—	—	—	—	—	—	1.341/0.000	—
11	1.259/0.420	1.340/0.002	—	—	—	—	—	—
12	1.250/0.440	1.340/0.001	—	—	—	—	—	—
13	1.259/0.420	—	—	—	—	—	—	—

applied at the central node (node 5), which is supported by the surrounding bars anchored at nodes located on the base plane. The structural configuration and nodal coordinates are reported in Table 4.9 and Table 4.10.




Figure 4.10: Eight-bar truss configuration used in Test Problem 2. The external load is applied at node 5.

Table 4.9: Element topology: node connectivity.

Element number	Node numbers
1	1-5
2	2-5
3	3-5
4	4-5
5	6-5
6	7-5
7	8-5
8	9-5

Table 4.10: Nodal coordinates (in mm).

Node	x	y	z
1	-250	-250	0
2	-250	250	0
3	250	250	0
4	250	-250	0
5	0	0	375
6	-375	0	0
7	0	375	0
8	375	0	0
9	0	-375	0

The applied load vector at node 5 is given by:

$$F_x = 40 \text{ kN} \quad F_y = 20 \text{ kN} \quad F_z = 200 \text{ kN}$$

The design variables are the cross-sectional areas $x_a \in \mathbb{R}_{>0}$ (expressed in mm^2) of the eight structural bars. Lower bounds $x_a \geq 100 \text{ mm}^2$ are imposed to avoid singular stiffness and ensure physical feasibility. No upper bounds are enforced, as the solution lies well within the feasible range. The optimization objective is the minimization of the total weight of the structure, subject to axial stress constraints on each bar:

$$\begin{aligned} \min_{x_a \geq 100} \quad & W(x_1, \dots, x_8) \\ \text{subject to} \quad & |\sigma_a(\mathbf{x})| \leq 100 \quad \text{for } a = 1, \dots, 8 \end{aligned} \quad (4.132)$$

where σ_a denotes the axial stress in element a , and $W(\mathbf{x})$ is the total structural weight, computed as the sum of the elemental contributions based on geometry and material density.

The initial design is uniformly set to $x_a^{(0)} = 400 \text{ mm}^2$. Move limits are defined according to a multiplicative rule to preserve relative variation:

$$\alpha_a^{(k)} = 0.5 x_a^{(k)} \quad \beta_a^{(k)} = 2.0 x_a^{(k)}$$

To prevent excessive curvature or instability in the convex approximation, the following bounds are imposed on the moving asymptotes:

$$-50 x_a^{(k)} \leq L_a^{(k)} \leq 0.4 x_a^{(k)} \quad 2.5 x_a^{(k)} \leq U_a^{(k)} \leq 50 x_a^{(k)}$$

During the first two iterations, the asymptotes are initialized as:

$$L_a^{(k)} = 0 \quad U_a^{(k)} = 5 x_a^{(k)}$$

From the third iteration onward, the adaptive update mechanism described in Section 4.3.3.2 is employed.

The optimal solution is reported in Table 4.11 and perfectly matches the results in [4], thus confirming the correctness of the implementation. The solution exhibits active stress constraints

Table 4.11: Optimal cross-sectional areas (in mm) obtained with MMA.

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8
880	720	260	520	100	100	100	100

only on the first four elements, while the remaining cross-section areas converge to the lower bound, indicating structural redundancy in those bars.

The influence of the moving asymptote update parameter s on the convergence behavior is assessed in Table 4.12, which reports the evolution of the structural mass over successive iterations. The results are compared with those obtained using a traditional formulation with fixed lower moving asymptotes $L_a^{(k)} = 0$. As clearly shown, larger values of s promote faster convergence in the early iterations by expanding the approximation domain and enabling more aggressive updates. Conversely, smaller values of s lead to slower but more stable progress due to tighter control over curvature.

The results further confirm the efficiency and robustness of MMA in constrained structural optimization. In particular, the adaptive update of the moving asymptotes plays a crucial role in accelerating convergence while maintaining numerical stability.

Table 4.12: Convergence of total mass (in kg) for Test Problem 2. Results match those reported in [4].

Iter.	Traditional	MMA ($s = \frac{3}{4}$)	MMA ($s = \frac{1}{2}$)	MMA ($s = \frac{1}{4}$)
0	13.05	13.05	13.05	13.05
1	11.68	12.10	12.10	12.10
2	11.66	11.67	11.67	11.67
3	11.64	11.65	11.65	11.65
4	11.62	11.64	11.63	11.61
5	11.60	11.62	11.60	11.52
6	11.59	11.60	11.53	11.42
7	11.57	11.56	11.44	11.28
8	11.55	11.52	11.35	11.23
9	11.53	11.47	11.25	–
10	11.52	11.41	11.23	–
11	11.50	11.36	–	–
12	11.48	11.31	–	–
13	11.46	11.24	–	–
14	11.45	11.23	–	–
15	11.43	–	–	–
⋮	⋮	⋮	⋮	⋮
30	11.27	–	–	–
⋮	⋮	⋮	⋮	⋮
39	11.23	–	–	–

4.3.4.3 Two-Bar Truss with Mixed Design Variables

The third benchmark problem concerns a planar truss composed of two identical bars forming a symmetric triangular configuration, as illustrated in Figure 4.11. The structure is supported at

nodes 1 and 2, and a vertical load is applied at the free apex node (node 3). The design variables are: (i) the cross-sectional area $x_1 \in \mathbb{R}_{>0}$ (in cm^2) shared by both members, and (ii) the half-span distance $x_2 \in \mathbb{R}_{>0}$ (in meters) between the two supports. The elements topology and nodal configuration are given in Tables 4.13 and 4.14, respectively.




Figure 4.11: Two-bar truss configuration (Test Problem 3). The structure is symmetric, with both bars having identical cross-sectional area x_1 , and base length defined by $2x_2$. An external load \mathbf{F} is applied at node 3.

Table 4.13: Element topology: node connectivity.

Element number	Node numbers
1	1-3
2	2-3

Table 4.14: Nodal coordinates (in meters).

Node	x	y
1	$-x_2$	0.0
2	x_2	0.0
3	0.0	1.0

The external load vector \mathbf{F} , applied at node 3, has the following components:

$$F_x = 24.8 \text{ kN}, \quad F_y = 198.4 \text{ kN} \quad |\mathbf{F}| = 200 \text{ kN}$$

The objective is to minimize the structural weight, subject to axial stress constraints on both elements. The optimization problem is formulated as:

$$\begin{aligned}
& \min_{x_1, x_2} w(x_1, x_2) = C_1 x_1 \sqrt{1 + x_2^2} \\
& \text{subject to } \sigma_1(x_1, x_2) = C_2 \sqrt{1 + x_2^2} \left(\frac{8}{x_1} + \frac{1}{x_1 x_2} \right) \leq 1 \\
& \sigma_2(x_1, x_2) = C_2 \sqrt{1 + x_2^2} \left(\frac{8}{x_1} - \frac{1}{x_1 x_2} \right) \leq 1
\end{aligned} \tag{4.133}$$

where $C_1 = 1.0$ is a normalization factor for the weight, and $C_2 = 0.124$ captures the stress scaling due to geometry and loading. The feasible domain is defined by box constraints:

$$0.2 \leq x_1 \leq 4.0 \quad (\text{cm}^2) \quad 0.1 \leq x_2 \leq 1.6 \quad (\text{m})$$

Note that the stress expressions in Equation (4.133) reflect the coupling between the cross-section area and the geometric layout. Due to the inverse dependence on both x_1 and x_2 , the constraint functions are nonlinear and nonconvex, posing a suitable challenge for the approximation capabilities of MMA.

The initial design is selected as $x_1^{(0)} = 1.50 \text{ cm}^2$, $x_2^{(0)} = 0.50 \text{ m}$, which corresponds to a feasible but suboptimal configuration. Table 4.15 reports the evolution of the design variables (x_1, x_2) , the dominant stress σ_1 , and the structural weight w over successive iterations of the MMA.

Table 4.15: Results for Test Problem 3 using MMA. Each row shows x_1 , x_2 , stress σ_1 , and weight w .

Iteration	x_1	x_2	σ_1	w
0	1.50	0.50	0.92	1.68
1	1.39	0.25	1.10	1.43
2	1.22	0.50	1.13	1.37
3	1.39	0.25	1.10	1.44
4	1.37	0.38	1.03	1.47
5	1.41	0.38	1.00	1.51

The results confirm that MMA converges in a small number of iterations to a feasible and nearly optimal configuration, despite the absence of explicit move limits. The stress constraints are satisfied at the final iterate, with the critical stress $\sigma_1 \approx 1.00$, indicating that the optimal design lies at the boundary of the admissible set. The evolution of the design variables suggests a transient oscillatory behavior between configurations with narrow and wide support spacing, which is progressively attenuated through the adaptive regulation of the asymptotes. This behavior highlights the sensitivity of the stress constraints to geometric coupling, as well as the stabilizing role played by the MMA curvature control mechanism.

The numerical results are in excellent agreement with those reported in [4], thereby validating the correctness of the implementation and confirming the effectiveness of MMA in solving non-linear constrained optimization problems involving simultaneous sizing and geometric variables.

4.4 Conclusions

The present chapter has investigated two established methodologies in structural optimization, namely the density-based SIMP approach combined with an Optimality Criteria (OC) scheme, and the Method of Moving Asymptotes (MMA). The analysis has focused on a rigorous and unified presentation of their theoretical foundations, together with a systematic assessment of their numerical behavior on representative benchmark problems. The objective has not been the introduction of new algorithmic developments, but rather the clarification, reinterpretation, and consistent formulation of methods that are widely used in computational mechanics.

Section 4.2 has addressed topology optimization through the SIMP framework in two-dimensional settings, with emphasis on compliance minimization under volume constraints. The formulation has been examined within a Lagrangian setting, highlighting the role of density interpolation, sensitivity analysis, and filtering techniques. The numerical results confirm that the combined SIMP–OC approach provides a robust and effective strategy for material distribution problems. In particular, density-based and sensitivity-based filtering have been analysed in detail, showing their fundamental role in controlling mesh dependency, improving numerical stability, and ensuring physically meaningful solutions. The agreement with reference benchmark results supports the correctness of the implementation and the consistency of the adopted formulation.

Section 4.3 has provided a comprehensive and self-contained analysis of the Method of Moving Asymptotes, with particular emphasis on its analytical structure and approximation mechanisms. The derivation of convex and separable local models has been presented starting from first principles, together with a detailed discussion of curvature control through moving asymptotes and adaptive update strategies. The formulation has been developed from a one-dimensional setting to multidimensional and constrained problems, thereby offering a unified interpretation of the method. The numerical experiments confirm the capability of MMA to handle nonlinear and constrained optimization problems in a stable and efficient manner, provided that appropriate parameter selection and update rules are adopted. The comparison with results available in the literature further validates the implementation.

The methodologies analysed in the present chapter address different classes of structural optimization problems and rely on distinct algorithmic principles. The SIMP approach, combined with OC updates, is particularly suited for large-scale topology optimization problems involving distributed design variables, whereas MMA provides a general framework for constrained nonlinear optimization, applicable to sizing, shape, and mixed-variable formulations. Although both approaches rely on local approximations and iterative updates, their mathematical structures, underlying assumptions, and domains of applicability remain fundamentally different.

From a broader perspective, the present chapter provides a coherent theoretical and computational foundation for the optimization techniques employed in structural mechanics. Such a foundation is essential for the integration of optimization methods with the uncertainty-aware frameworks developed in Chapters 2 and 3. The formulation of both SIMP-based topology optimization and MMA within a deterministic setting enables a clear identification of the algorithmic components that must be extended in order to account for uncertainty.

In particular, the extension of topology optimization to uncertainty-aware settings requires the incorporation of stochastic representations of material properties, loading conditions, and model parameters, together with the introduction of robust or reliability-based formulations. Similarly, the integration of MMA within surrogate-assisted Bayesian frameworks offers a natural pathway toward inverse design under uncertainty, where optimization and probabilistic inference are combined in a consistent manner. The analytical structure of MMA, based on convex separable approximations and explicit curvature control, is particularly suitable for such extensions.

Future research directions therefore include the development of uncertainty-aware topology optimization strategies, the coupling of MMA with surrogate modeling and Bayesian inversion techniques, and the extension of the presented formulations to large-scale and multi-physics problems. Additional efforts are also required to establish systematic procedures for parameter selection, adaptive approximation strategies, and convergence analysis in the presence of uncertainty. These perspectives highlight the role of the methodologies analysed in this chapter as a foundational component for the development of advanced design frameworks integrating optimization, surrogate modeling, and uncertainty quantification.

Chapter 5

Conclusion and Future perspectives

The present doctoral thesis has investigated methodologies for modeling, inverse analysis, and design in engineering systems characterized by high-dimensional parameter spaces, nonlinear behavior, and multiple sources of uncertainty. The work is structured around three complementary but distinct research directions addressing key challenges in computational mechanics: uncertainty quantification for thermomechanical modeling in metal additive manufacturing, surrogate-based Bayesian inversion for inverse design of architected materials, and structural topology optimization.

These research directions have been developed in different application contexts and rely on distinct methodological approaches. They do not constitute a unified framework; rather, they provide complementary perspectives on how surrogate modelling, probabilistic inference, and optimization techniques can be employed to address complex engineering problems under limited data and significant uncertainty. At the same time, they share common methodological foundations that enable a meaningful comparison and suggest possible directions for integration.

Taken together, the three contributions reflect a progressive methodological development, moving from forward uncertainty quantification, to inverse inference, and ultimately toward design-oriented methodologies. This progression highlights the role of surrogate modeling and probabilistic reasoning as unifying elements that support the transition from analysis to decision-making in computational engineering.

The first contribution develops an uncertainty-quantification framework for thermomechanical modeling in metal additive manufacturing. The approach combines variance-based global sensitivity analysis, Bayesian inverse calibration, and forward uncertainty propagation. High-dimensional responses are approximated using single- and multi-fidelity sparse-grid surrogate models, enabling efficient exploration of the parameter space and probabilistic prediction of quantities of interest. The resulting framework is particularly suited for forward uncertainty propagation and sensitivity analysis in settings where the response exhibits sufficient regularity and where structured approximations of the parameter space are effective.

The second contribution addresses inverse problems and inverse design through surrogate-based Bayesian methodologies. Gaussian-process surrogate models are constructed and adaptively refined based on predictive uncertainty, providing flexible approximations of computationally expensive forward models together with a local measure of predictive confidence. Within this setting, the inverse problem is formulated in a probabilistic framework and solved using approximate posterior representations based on local Laplace approximations. The formulation explicitly accounts for surrogate-induced uncertainty and residual discrepancy, and is applied to

the inverse design of spinodoid architected materials, where the underlying structure–property relationships are nonlinear and sensitive to parameter variations.

A direct comparison between the surrogate modeling strategies adopted in Chapters 2 and 3 highlights their complementary roles. Sparse-grid interpolation provides an efficient and structured representation when the input space has moderate dimensionality and the model response is sufficiently smooth, making it well suited for forward uncertainty propagation, global sensitivity analysis, and scenarios requiring a global approximation of the parameter space. In contrast, Gaussian-process regression offers greater flexibility in representing complex, nonlinear, and irregular responses, together with an intrinsic quantification of predictive uncertainty that supports adaptive sampling and targeted refinement. These features make Gaussian-process models particularly effective in inverse problems and design settings, where computational resources must be concentrated in regions of high posterior probability or near optimal configurations. The choice between these approaches therefore depends on the interplay between dimensionality, smoothness of the response, and the specific objectives of the analysis, namely global exploration versus local refinement.

The third contribution examines structural topology and nonlinear optimization methods, with particular emphasis on the Optimality Criterion method within the Solid Isotropic Material with Penalization framework and on the Method of Moving Asymptotes. The analysis provides a rigorous clarification of the underlying theoretical structure of these established methodologies, including the role of local convex approximations, curvature control, and constraint handling in ensuring numerical stability and convergence. The study is conducted in a deterministic setting and does not introduce new optimization algorithms, but rather aims at consolidating and interpreting existing methods.

The absence of uncertainty representation in the optimization framework highlights a natural direction for further developments. The methodologies introduced in Chapters 2 and 3 provide a basis for extending structural optimization toward uncertainty-aware formulations. In particular, sparse-grid surrogate models may be employed to approximate the response of parametrized structures and enable efficient forward uncertainty propagation and sensitivity analysis within topology optimization problems. Similarly, Gaussian-process surrogate models may be used to construct local probabilistic approximations of structural responses, supporting adaptive sampling strategies and refinement in regions characterized by high design sensitivity or uncertainty. Such extensions would enable the formulation of robust and reliability-based topology optimization problems accounting for variability in material properties, loading conditions, and modeling assumptions.

More generally, the three research directions suggest the possibility of defining a coherent computational workflow integrating forward uncertainty quantification, Bayesian inversion, and structural optimization. Within such a workflow, surrogate models would approximate computationally expensive forward models, Bayesian inversion would incorporate observational data and quantify parameter uncertainty, and optimization techniques would exploit this information to identify designs that satisfy prescribed performance criteria under uncertainty. This integration would enable a consistent propagation of uncertainty from model inputs to design variables, bridging the gap between analysis and design. Such a framework is not developed in the present

work and represents a natural direction for future research.

Future developments may further include the incorporation of measurement data and observational uncertainty within inverse design formulations, the adoption of sampling-based Bayesian inference methods to improve the characterization of posterior distributions beyond local approximations, and the extension to higher-dimensional parameterizations through advanced surrogate modeling strategies, including adaptive multi-fidelity approaches.

In conclusion, the thesis provides three complementary methodological contributions addressing uncertainty quantification, inverse analysis, and structural optimization. Although developed independently, these contributions share consistent methodological principles related to surrogate modeling, probabilistic inference, and computational efficiency. These common elements establish a foundation for future developments aimed at the systematic integration of uncertainty quantification, Bayesian inference, and optimization in the analysis and design of complex engineering systems.

Bibliography

- [1] National Institute of Standards and Technologies (NIST) - 2018 AM-Bench Test Descriptions for AMB2018-01, 2018.
- [2] Lyle Levine, Brandon Lane, Jarred Heigel, Kalman Migler, Mark Stoudt, Thien Phan, Richard Ricker, Maria Strantza, Michael Hill, Fan Zhang, et al. Outcomes and conclusions from the 2018 AM-bench measurements, challenge problems, modeling submissions, and conference. *Integrating Materials and Manufacturing Innovation*, 9:1–15, 2020.
- [3] S Ivvan Valdez, Salvador Botello, Miguel A Ochoa, José L Marroquín, and Victor Cardoso. Topology optimization benchmarks in 2d: Results for minimum compliance and minimum volume in planar stress problems. *Archives of Computational Methods in Engineering*, 24: 803–839, 2017.
- [4] Krister Svanberg. The method of moving asymptotes — a new method for structural optimization. *International journal for numerical methods in engineering*, 24(2):359–373, 1987.
- [5] Roger Ghanem, David Higdon, and Houman Owhadi. *Handbook of uncertainty quantification*, volume 6. Springer, 2017.
- [6] Christian Soize. *Uncertainty Quantification*. Springer, 2017.
- [7] Ralph C Smith. *Uncertainty quantification: Theory, Implementation, and Applications*, volume 12. SIAM, 2013.
- [8] Timothy John Sullivan. *Introduction to uncertainty quantification*, volume 63. Springer, 2015.
- [9] Tarasankar DebRoy, Huiliang L Wei, James S Zuback, Tuhin Mukherjee, John W Elmer, John O Milewski, Allison Michelle Beese, Alexander Wilson-Heid, Amitava De, and Wei Zhang. Additive manufacturing of metallic components—process, structure and properties. *Progress in materials science*, 92:112–224, 2018.
- [10] Julia R Greer and Vikram S Deshpande. Three-dimensional architected materials and structures: Design, fabrication, and mechanical behavior. *MRS Bulletin*, 44(10):750–757, 2019.

- [11] George IN Rozvany and Tomasz Lewiński. Topology optimization in structural and continuum mechanics. 2014.
- [12] Ian Gibson, David Rosen, Brent Stucker, Mahyar Khorasani, David Rosen, Brent Stucker, and Mahyar Khorasani. *Additive Manufacturing Technologies*, volume 17. Springer, 2021.
- [13] William E Frazier. Metal additive manufacturing: a review. *Journal of Materials Engineering and performance*, 23(6):1917–1928, 2014.
- [14] Wayne E. King, Andrew T. Anderson, Robert M. Ferencz, Neil E. Hodge, Chandrika Kamath, Saad A. Khairallah, and Alexander M. Rubenchik. Laser powder bed fusion additive manufacturing of metals; physics, computational, and materials challenges. *Applied Physics Reviews*, 2(4):041304, 2015.
- [15] Michele Chiumenti, Eric Neiva, Emilio Salsi, Miguel Cervera, Santiago Badia, Joan Moya, Zhuoer Chen, Caroline Lee, and Christopher Davies. Numerical modelling and experimental validation in Selective Laser Melting. *Additive Manufacturing*, 18:171–185, 2017.
- [16] Erik R. Denlinger, Michael Gouge, Jeff Irwin, and Pan Michaleris. Thermomechanical Model Development and in situ Experimental Validation of the Laser Powder-Bed Fusion Process. *Additive Manufacturing*, 16:73–80, 2017.
- [17] Yang Liu, Jing Shi, and Yachao Wang. Evolution, control, and mitigation of residual stresses in additively manufactured metallic materials: a review. *Advanced Engineering Materials*, 25(22):2300489, 2023.
- [18] Nabin Bastola, Muhammad P Jahan, Nithin Rangasamy, and Chandra Sekhar Rakurty. A review of the residual stress generation in metal additive manufacturing: analysis of cause, measurement, effects, and prevention. *Micromachines*, 14(7):1480, 2023.
- [19] Katerina Giannoukou, Stefano Marelli, and Bruno Sudret. Uncertainty-aware multifidelity surrogate modeling with noisy data. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 11(3):04025037, 2025.
- [20] Zeynab Azarhoosh and Majid Ilchi Ghazaan. A review of recent advances in surrogate models for uncertainty quantification of high-dimensional engineering applications. *Computer Methods in Applied Mechanics and Engineering*, 433:117508, 2025.
- [21] Siddhant Kumar, Stephanie Tan, Li Zheng, and Dennis M Kochmann. Inverse-designed spinodoid metamaterials. *npj Computational Materials*, 6(1):73, 2020.
- [22] Li Zheng, Siddhant Kumar, and Dennis M Kochmann. Data-driven topology optimization of spinodoid metamaterials with seamlessly tunable anisotropy. *Computer Methods in Applied Mechanics and Engineering*, 383:113894, 2021.

- [23] Alexandra Otto, Max Rosenkranz, Karl A Kalina, and Markus Kästner. Data-driven inverse design of spinodoid architected materials. *GAMM-Mitteilungen*, 48(4):e70008, 2025.
- [24] Mahtab Vafaeefer, Kevin M Moerman, Majid Kavousi, and Ted J Vaughan. A morphological, topological and mechanical investigation of gyroid, spinodoid and dual-lattice algorithms as structural models of trabecular bone. *Journal of the Mechanical Behavior of Biomedical Materials*, 138:105584, 2023.
- [25] Marc GD Geers, VG Kouznetsova, and WAM Brekelmans. Computational homogenization of heterogeneous materials. *Journal of Computational and Applied Mathematics*, 234(7): 2175–2182, 2010.
- [26] Justin Dirrenberger, Samuel Forest, and Dominique Jeulin. Computational homogenization of architected materials. In *Architected Materials in Nature and Engineering: Archimats*, pages 89–139. Springer, 2019.
- [27] Reza Alizadeh, Janet K Allen, and Farrokh Mistree. Managing computational complexity using surrogate models: a critical review. *Research in Engineering Design*, 31(3):275–298, 2020.
- [28] Jakub Kudela and Radomil Matousek. Recent advances and applications of surrogate models for finite element method computations: a review: J. kudela, r. matousek. *Soft Computing*, 26(24):13709–13733, 2022.
- [29] Jerome Sacks, William J Welch, Toby J Mitchell, and Henry P Wynn. Design and analysis of computer experiments. *Statistical science*, 4(4):409–423, 1989.
- [30] Marc C Kennedy and Anthony O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.
- [31] Martin Philip Bendsoe and Ole Sigmund. *Topology optimization: theory, methods, and applications*. Springer Science & Business Media, 2003.
- [32] Ole Sigmund and Kurt Maute. Topology optimization approaches: A comparative review. *Structural and Multidisciplinary Optimization*, 48(6):1031–1055, 2013.
- [33] K Svanberg. The method of moving asymptotes (MMA) with some extensions. In *Optimization of large structural systems*, pages 555–566. Springer, 1993.
- [34] Erik Andreassen, Anders Clausen, Mattias Schevenels, Boyan S Lazarov, and Ole Sigmund. Efficient topology optimization in MATLAB using 88 lines of code. *Structural and Multidisciplinary Optimization*, 43(1):1–16, 2011.

- [35] Ole Sigmund and Joakim Petersson. Numerical instabilities in topology optimization: a survey on procedures dealing with checkerboards, mesh-dependencies and local minima. *Structural optimization*, 16(1):68–75, 1998.
- [36] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. SIAM, 2005.
- [37] Andrew M Stuart. Inverse problems: a bayesian perspective. *Acta numerica*, 19:451–559, 2010.
- [38] Daniela Calvetti and Erkki Somersalo. Inverse problems: From regularization to bayesian inference. *Wiley Interdisciplinary Reviews: Computational Statistics*, 10(3):e1427, 2018.
- [39] Kaufui V. Wong and Aldo Hernandez. A review of additive manufacturing. *International scholarly research notices*, 2012:10 pages, 2012.
- [40] Sohini Chowdhury, N. Yadaiah, Chander Prakash, Seeram Ramakrishna, Saurav Dixit, Lovi Raj Gulta, and Dharam Buddhi. Laser Powder Bed fusion: A State-of-the-Art Review of the Technology, Materials, Properties & Defects, and Numerical Modelling. *Journal of Materials Research and Technology*, 20:2109–2172, 2022.
- [41] Jacob Smith, Wei Xiong, Wentao Yan, Stephen Lin, Puikai Cheng, Orion L Kafka, Gregory J Wagner, Jian Cao, and Wing Kam Liu. Linking process, structure, property, and performance for metal-based additive manufacturing: computational approaches with experimental support. *Computational Mechanics*, 57(4):583–610, 2016.
- [42] Wayne King, Andrew T. Anderson, Robert M. Ferencz, Neil E. Hodge, Chandrika Kamath, and Saad A. Khairallah. Overview of Modelling and Simulation of Metal Powder Bed Fusion Process at Lawrence Livermore National Laboratory. *Materials Science and Technology*, 31(8):957–968, 2015.
- [43] Trevor Keller, Greta Lindwall, Supriyo Ghosh, Li Ma, Brandon M. Lane, Fan Zhang, Ursula R. Kattner, Eric A. Lass, Jarred C. Heigel, Yaakov Idell, et al. Application of finite element, phase-field, and CALPHAD-based methods to additive manufacturing of Ni-based superalloys. *Acta materialia*, 139:244–253, 2017.
- [44] Mohamad Bayat, Wen Dong, Jesper Thorborg, Albert C. To, and Jesper H. Hattel. A Review of Multi-scale and Multi-physics Simulations of Metal Additive Manufacturing Processes with Focus on Modeling Strategies. *Additive Manufacturing*, 47:102278, 2021.
- [45] Kubra Karayagiz, Luke Johnson, Raiyan Seede, Vahid Attari, Bing Zhang, Xueqin Huang, Supriyo Ghosh, Thien Duong, Ibrahim Karaman, Alaa Elwany, and Raymundo Arróyave. Finite interface dissipation phase field modeling of Ni–Nb under additive manufacturing conditions. *Acta Materialia*, 185:320–339, 2020.

- [46] Supriyo Ghosh, Kevin McReynolds, Jonathan E. Guyer, and Dilip Banerjee. Simulation of temperature, stress and microstructure fields during laser deposition of Ti–6Al–4V. *Modelling and simulation in materials science and engineering*, 26(7):075005, 2018.
- [47] Liang-Xing Lu, N. Sridhar, and Yong-Wei Zhang. Phase Field Simulation of Powder Bed-based Additive Manufacturing. *Acta Materialia*, 144:801–809, 2018.
- [48] A. Zinoviev, O. Zinovieva, V. Ploshikhin, V. Romanova, and R. Balokhonov. Evolution of Grain Structure During Laser Additive Manufacturing. Simulation by a Cellular Automata Method. *Materials & Design*, 106:321–329, 2016.
- [49] Elham Attar and Carolin Körner. Lattice Boltzmann Model for Thermal free Surface Flows with Liquid–Solid Phase Transition. *International Journal of Heat and Fluid Flow*, 32(1):156–163, 2011.
- [50] Alexander Klassen, Thorsten Scharowsky, and Carolin Körner. Evaporation model for beam based additive manufacturing using free surface lattice Boltzmann methods. *Journal of Physics D: Applied Physics*, 47(27):275303, 2014.
- [51] Paul A. Hooper. Melt Pool Temperature and Cooling Rates in Laser Powder Bed Fusion. *Additive Manufacturing*, 22:548–559, 2018.
- [52] John C. Steuben, Athanasios P. Iliopoulos, and John G. Michopoulos. Discrete element modeling of particle-based additive manufacturing processes. *Computer Methods in Applied Mechanics and Engineering*, 305:537–561, 2016.
- [53] Massimo Carraturo, John Jomo, Stefan Kollmannsberger, Alessandro Reali, Ferdinando Auricchio, and Ernst Rank. Modeling and Experimental Validation of an Immersed Thermo-mechanical Part-scale Analysis for Laser Powder Bed Fusion Processes. *Additive Manufacturing*, 36:101498, 2020.
- [54] Yingli Li, Kun Zhou, Pengfei Tan, Shu Beng Tor, Chee Kai Chua, and Kah Fai Leong. Modeling temperature and residual stress fields in selective laser melting. *International Journal of Mechanical Sciences*, 136:24–35, 2018.
- [55] Pengfei Tan, Fei Shen, Biao Li, and Kun Zhou. A thermo-metallurgical-mechanical model for selective laser melting of ti6al4v. *Materials & Design*, 168:107642, 2019.
- [56] Massimo Carraturo, Alex Viguerie, Alessandro Reali, and Ferdinando Auricchio. Two-level method part-scale thermal analysis of laser powder bed fusion additive manufacturing. *Engineering with Computers*, pages 1–14, 2022.
- [57] Alex Viguerie, Massimo Carraturo, Alessandro Reali, and Ferdinando Auricchio. A spatiotemporal two-level method for high-fidelity thermal analysis of laser powder bed fusion. *Finite Elements in Analysis and Design*, 210:103815, 2022.

- [58] Zhen Hu and Sankaran Mahadevan. Uncertainty quantification in prediction of material properties during additive manufacturing. *Scripta materialia*, 135:135–140, 2017.
- [59] Zhen Hu and Sankaran Mahadevan. Uncertainty Quantification and Management in Additive Manufacturing: Current Status, Needs, and Opportunities. *The International Journal of Advanced Manufacturing Technology*, 93:2855–2874, 2017.
- [60] James O. Berger. *Statistical decision theory and Bayesian analysis*. Springer Science & Business Media, 2013.
- [61] A. M. Stuart. Inverse Problems: A Bayesian Perspective. *Acta numerica*, 19:451–559, 2010.
- [62] T. Bui-Thanh, O. Ghattas, J. Martin, and G. Stadler. A Computational Framework for Infinite-Dimensional Bayesian Inverse Problems Part I: The Linearized Case, with Application to Global Seismic Inversion. *SIAM Journal on Scientific Computing*, 35(6): A2494–A2523, 2013.
- [63] Chandrika Kamath. Data mining and statistical inference in selective laser melting. *The International Journal of Advanced Manufacturing Technology*, 86:1659–1677, 2016.
- [64] A. Garg, K. Tai, and M.M. Savalani. State-of-the-art in empirical modelling of rapid prototyping processes. *Rapid Prototyping Journal*, 20(2):164–178, 2014.
- [65] Felipe Lopez, Paul Witherell, and Brandon Lane. Identifying Uncertainty in Laser Powder Bed Fusion Additive Manufacturing Models. *Journal of Mechanical Design*, 138(11): 114502, 2016.
- [66] Marco Grasso and Bianca Maria Colosimo. Process Defects and in Situ Monitoring Methods in Metal Powder Bed Fusion: A Review. *Measurement Science and Technology*, 28(4):044005, 2017.
- [67] Nina Korshunova, Iason Papaioannou, Stefan Kollmannsberger, Daninel Straub, and Ernst Rank. Uncertainty Quantification of Microstructure variability and Mechanical Behavior of Additively Manufactured Lattice Structures. *Computer Methods in Applied Mechanics and Engineering*, 385:114049, 2021.
- [68] P. Nath, Z. Hu, and S. Mahadevan. Multi-Level Uncertainty Quantification in Additive Manufacturing. In *2017 International Solid Freeform Fabrication Symposium*. University of Texas at Austin, 2017.
- [69] Ziyu Xie, Wen Jiang, Congjian Wang, and Xu Wu. Bayesian Inverse Uncertainty Quantification of a MOOSE-based Melt Pool Model for Additive Manufacturing Using Experimental Data. *Annals of Nuclear Energy*, 165:108782, 2022.

- [70] AmirPouya Hemmasian, Francis Ogoke, Parand Akbari, Jonathan Malen, Jack Beuth, and Amir Barati Farimani. Surrogate modeling of melt pool temperature field using deep learning. *Additive Manufacturing Letters*, 5:100123, 2023.
- [71] Matthias Speich, Carsten F Dormann, and Florian Hartig. Sequential monte-carlo algorithms for bayesian model calibration—a review and method comparison. *Ecological Modelling*, 455:109608, 2021.
- [72] Zhuo Wang, Chen Jiang, Pengwei Liu, Wenhua Yang, Ying Zhao, Mark F. Horstemeyer, Long-Qing Chen, Zhen Hu, and Lei Chen. Uncertainty Quantification and Reduction in Metal Additive Manufacturing. *Npj Computational Materials*, 6(1):175, 2020.
- [73] Paromita Nath, Zhen Hu, and Sankaran Mahadevan. Modeling and Uncertainty Quantification of Material Properties in Additive Manufacturing. In *2018 AIAA Non-Deterministic Approaches Conference*. 2018 AIAA Non-Deterministic Approaches Conference, 2018.
- [74] Pejman Honarmandi and Raymundo Arróyave. Uncertainty quantification and propagation in computational materials science and simulation-assisted materials design. *Integrating Materials and Manufacturing Innovation*, 9(1):103–143, 2020.
- [75] Supriyo Ghosh, Mohamad Mahmoudi, Luke Johnson, Alaa Elwany, Raymundo Arroyave, and Douglas Allaire. Uncertainty analysis of microsegregation during laser powder bed fusion. *Modelling and Simulation in Materials Science and Engineering*, 27(3):034002, 2019.
- [76] Mihaela Chiappetta, Chiara Piazzola, Massimo Carraturo, Lorenzo Tamellini, Alessandro Reali, and Ferdinando Auricchio. Sparse-grids uncertainty quantification of part-scale additive manufacturing processes. *International Journal of Mechanical Sciences*, 256:108476, 2023.
- [77] Mihaela Chiappetta, Chiara Piazzola, Lorenzo Tamellini, Alessandro Reali, Ferdinando Auricchio, and Massimo Carraturo. Data-informed uncertainty quantification for laser-based powder bed fusion additive manufacturing. *International Journal for Numerical Methods in Engineering*, 125(17):e7542, 2024.
- [78] H.J Bungartz and M. Griebel. Sparse grids. *Acta Numer.*, 13:147–269, 2004. ISSN 0962-4929.
- [79] Fabio Nobile, Raul Tempone, and Clayton Webster. Sparse grid stochastic collocation for elliptic pdes with random input data. *SIAM Journal on Numerical Analysis*, 2008.
- [80] Chiara Piazzola and Lorenzo Tamellini. The Sparse Grids Matlab kit—a Matlab Implementation of Sparse Grids for High-Dimensional Function Approximation and Uncertainty Quantification. *arXiv preprint arXiv:2203.09314*, 2022.

- [81] Joakim Beck, Lorenzo Tamellini, and Raúl Tempone. IGA-based Multi-Index Stochastic Collocation for random PDEs on arbitrary domains. *Computer Methods in Applied Mechanics and Engineering*, 351:330–350, 2019.
- [82] Abdul-Lateef Haji-Ali, Fabio Nobile, Lorenzo Tamellini, and Raúl Tempone. Multi-index stochastic collocation for random PDEs. *Computer Methods in Applied Mechanics and Engineering*, 306:95–122, 2016.
- [83] Massimo Carraturo, Brandon Lane, Ho Yeung, Stefan Kollmannsberger, Alessandro Reali, and Ferdinando Auricchio. Numerical Evaluation of Advanced Laser Control Strategies Influence on Residual Stresses for Laser Powder Bed Fusion Systems. *Integrating Materials and Manufacturing Innovation*, 9:435–445, 2020.
- [84] Babis Schoinochoritis, Dimitrios Chantzis, and Konstantinos Salonitis. Simulation of Metallic Powder Bed Additive Manufacturing Processes with the Finite Element Method: A Critical Review. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, 231(1):96–117, 2017.
- [85] Louis C Burmeister. *Convective heat transfer*, volume 2. John Wiley & Sons, 1993.
- [86] Akram Chergui, Nicolas Beraud, Frédéric Vignat, and François Villeneuve. Finite Element Modeling and Validation of Metal Deposition in Wire Arc Additive Manufacturing. In *International Joint Conference on Mechanics, Design Engineering & Advanced Manufacturing*. Springer, 2020.
- [87] Ke An, Lang Yuan, Laura Dial, Ian Spinelli, Alexandru D. Stoica, and Yan Gao. Neutron Residual Stress Measurement and Numerical Modeling in a Curved Thin-walled Structure by Laser Powder Bed Fusion Additive manufacturing. *Materials & design*, 135:122–132, 2017.
- [88] Zhengtao Gan, Yanping Lian, Stephen E. Lin, Kevontrez K. Jones, Wing Kam Liu, and Gregory J. Wagner. Benchmark Study of Thermal Behavior, Surface Topography, and Dendritic Microstructure in Selective Laser Melting of Inconel 625. *Integrating Materials and Manufacturing Innovation*, 8:178–193, 2019.
- [89] Yiğit M. Arsoy, Luis E. Criales, and Tuğrul Özel. Modeling and Simulation of Thermal Field and Solidification in Laser Powder Bed Fusion of Nickel Alloy IN625. *Optics & Laser Technology*, 109:278–292, 2019.
- [90] Chao Li, Erik R. Denlinger, Michael F. Gouge, Jeff E. Irwin, and Pan Michaleris. Numerical Verification of an Octree Mesh Coarsening Strategy for Simulating Additive Manufacturing Processes. *Additive Manufacturing*, 30:100903, 2019.
- [91] Kyung-Min Hong, Corbin M. Grohol, and Yung C. Shin. Comparative Assessment of Physics-based Computational Models on the NIST Benchmark Study of Molten Pool

- Dimensions and Microstructure for Selective Laser Melting of Inconel 625. *Integrating Materials and Manufacturing Innovation*, 10:58–71, 2021.
- [92] I. Babuška, F. Nobile, and R. Tempone. A Stochastic Collocation Method for Elliptic Partial Differential Equations with Random Input Data. *SIAM Review*, 52(2):317–355, 2010.
- [93] Dongbin Xiu and Jan S. Hesthaven. High-order Collocation Methods for Differential Equations with Random Inputs. *SIAM Journal on Scientific Computing*, 27(3):1118–1139, 2005.
- [94] Hans-Joachim Bungartz and Michael Griebel. Sparse Grids. *Acta numerica*, 13:147–269, 2004.
- [95] Alfio Quarteroni, Riccardo Sacco, and Fausto Saleri. *Numerical mathematics*, volume 37. Springer Science & Business Media, 2010.
- [96] Ilya M. Sobol. Global Sensitivity Indices for Nonlinear Mathematical Models and Their Monte Carlo Estimates. *Mathematics and computers in simulation*, 55(1-3):271–280, 2001.
- [97] Andrea Saltelli, Stefano Tarantola, Francesca Campolongo, Marco Ratto, et al. *Sensitivity analysis in practice: a guide to assessing scientific models*, volume 1. Wiley Online Library, 2004.
- [98] Andrea Saltelli, Marco Ratto, Terry Andres, Francesca Campolongo, Jessica Cariboni, Debora Gatelli, Michaela Saisana, and Stefano Tarantola. *Global sensitivity analysis: the primer*. John Wiley & Sons, 2008.
- [99] Bruno Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability engineering & system safety*, 93(7):964–979, 2008.
- [100] Bertrand Iooss and Paul Lemaître. A review on global sensitivity analysis methods. *Uncertainty management in simulation-optimization of complex systems: algorithms and applications*, 59:101–122, 2015.
- [101] Bruno Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, 93(7):964–979, 2008.
- [102] Thomas Crestaux, Olivier Le Maître, and Jean-Marc Martinez. Polynomial chaos expansion for sensitivity analysis. *Reliability Engineering & System Safety*, 94(7):1161–1172, 2009.
- [103] Loïc Le Gratiet, Stefano Marelli, and Bruno Sudret. *Uncertainty Quantification with Polynomial Chaos Expansions*. Springer, 2017.

- [104] Stefano Marelli and Bruno Sudret. Uqlab: A framework for uncertainty quantification in matlab. *Proceedings of the 2nd International Conference on Vulnerability and Risk Analysis and Management*, 2014.
- [105] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of Markov Chain Monte Carlo*. CRC press, 2011.
- [106] Paul Marjoram, John Molitor, Vincent Plagnol, and Simon Tavaré. Markov Chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 100(26): 15324–15328, 2003.
- [107] J-L Gauvain and Chin-Hui Lee. Maximum a posteriori estimation for multivariate gaussian mixture observations of markov chains. *IEEE transactions on speech and audio processing*, 2(2):291–298, 1994.
- [108] Robert Bassett and Julio Deride. Maximum a posteriori estimators as a limit of bayes estimators. *Mathematical Programming*, 174:129–144, 2019.
- [109] Emanuel Parzen. On estimation of a Probability Density Function and Mode. *The annals of Mathematical Statistics*, 33(3):1065–1076.
- [110] Murray Rosenblatt. Remarks on some nonparametric estimates of a density function. *The annals of Mathematical Statistics*, 27(3):832–837.
- [111] Håvard Rue, Sara Martino, and Nicolas Chopin. Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 71(2):319–392, 2009.
- [112] Sarat C Dass, Jaeyong Lee, Kyoungjae Lee, and Jonghun Park. Laplace based approximate posterior inference for differential equation models. *Statistics and Computing*, 27(3): 679–698, 2017.
- [113] A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, and J. Timmer. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25(15):1923–1929, 2009.
- [114] Paul Lartaud, Philippe Humbert, and Josselin Garnier. Sequential design for surrogate modeling in bayesian inverse problems. *arXiv preprint arXiv:2402.16520*, 2024.
- [115] John D. Jakeman, Michael S. Eldred, Gianluca Geraci, and Alex Gorodetsky. Adaptive multi-index collocation for uncertainty quantification and sensitivity analysis. *International Journal for Numerical Methods in Engineering*, 121(6):1314–1343, 2020.
- [116] Chiara Piazzola, Lorenzo Tamellini, Riccardo Pellegrini, Riccardo Brogna, Andrea Serani, and Matteo Diez. Comparing Multi-Index Stochastic Collocation and Multi-Fidelity Stochastic Radial Basis Functions for Forward Uncertainty Quantification of Ship Resistance. *Engineering with Computers*, 2022:1–29, 2022.

- [117] Jingyi Wang and Panayiotis Papadopoulos. Coupled thermomechanical analysis of fused deposition using the finite element method. *Finite Elements in Analysis and Design*, 197:103607, 2021.
- [118] M Giselle Fernández-Godino, Chanyoung Park, Nam-Ho Kim, and Raphael T Haftka. Review of multi-fidelity models. 2016. arXiv preprint arXiv:1609.07196.
- [119] B. Peherstorfer, K. Willcox, and M. Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review*, 60(3):550–591, 2018.
- [120] K.-T. Kim, U. Villa, M. Parno, Y. Marzouk, O. Ghattas, and N. Petra. HIPPYlib-MUQ: A Bayesian Inference Software Framework for Integration of Data with Complex Predictive Models under Uncertainty. *ACM Trans. Math. Softw.*, 49(2), 2023.
- [121] C. Piazzola, L. Tamellini, and R. Tempone. A note on tools for prediction under uncertainty and identifiability of SIR-like dynamical systems for epidemiology. *Mathematical Biosciences*, 332:108514, 2021.
- [122] Mihaela Chiappetta, Chiara Piazzola, Massimo Carraturo, Lorenzo Tamellini, Alessandro Reali, and Ferdinando Auricchio. Inverse and forward sparse-grids-based uncertainty quantification analysis of laser-based powder bed fusion of metals. *International Journal of Mechanical Sciences*, 256:108476, 2023.
- [123] Akil Narayan and John D. Jakeman. Adaptive Leja Sparse Grid Constructions for Stochastic Collocation and High-Dimensional Approximation. *SIAM Journal on Scientific Computing*, 36(6):A2952–A2983, 2014.
- [124] Fabio Nobile, Lorenzo Tamellini, and Raul Tempone. Comparison of Clenshaw–Curtis and Leja quasi-optimal sparse grids for the approximation of random PDEs. In *Spectral and High Order Methods for Partial Differential Equations ICOSAHOM 2014: Selected papers from the ICOSAHOM conference, June 23-27, 2014, Salt Lake City, Utah, USA*. Springer, 2015.
- [125] F. Nobile, L. Tamellini, and R. Tempone. Comparison of Clenshaw–Curtis and Leja Quasi-Optimal Sparse Grids for the Approximation of Random PDEs. In R. M. Kirby, M. Berzins, and J. S. Hesthaven, editors, *Spectral and High Order Methods for Partial Differential Equations - ICOSAHOM '14*, volume 106 of *Lecture Notes in Computational Science and Engineering*, pages 475–482. Springer International Publishing, 2015.
- [126] Sean Molesky, Zin Lin, Alexander Y Piggott, Weiliang Jin, Jelena Vucković, and Alejandro W Rodriguez. Inverse design in nanophotonics. *Nature Photonics*, 12(11):659–670, 2018.
- [127] Simon Arridge, Peter Maass, Ozan Öktem, and Carola-Bibiane Schönlieb. Solving inverse problems using data-driven models. *Acta numerica*, 28:1–174, 2019.

- [128] Richard C Aster, Brian Borchers, and Clifford H Thurber. *Parameter estimation and inverse problems*. Elsevier, 2018.
- [129] Andreas Kirsch et al. *An introduction to the mathematical theory of inverse problems*, volume 120. Springer, 2011.
- [130] Francisco Duarte Moura Neto and Antônio José da Silva Neto. *An introduction to inverse problems with applications*. Elsevier, 2018.
- [131] Steven N Evans and Philip B Stark. Inverse problems as statistics. *Inverse problems*, 18(4):R55–R97, 2002.
- [132] Sujin Bureerat and Nantiwat Pholdee. Inverse problem based differential evolution for efficient structural health monitoring of trusses. *Applied Soft Computing*, 66:462–472, 2018.
- [133] Michael S Zhdanov. *Inverse theory and applications in geophysics*, volume 36. Elsevier, 2015.
- [134] Alex Zunger. Inverse design in search of materials with target functionalities. *Nature Reviews Chemistry*, 2(4):0121, 2018.
- [135] Jia Wang, Yingxue Wang, and Yanan Chen. Inverse design of materials by machine learning. *Materials*, 15(5):1811, 2022.
- [136] Wei Liu, Ruoyu You, Jie Zhang, and Qingyan Chen. Development of a fast fluid dynamics-based adjoint method for the inverse design of indoor environments. *Journal of Building Performance Simulation*, 10(3):326–343, 2017.
- [137] Kelsey Allen, Tatiana Lopez-Guevara, Kimberly L Stachenfeld, Alvaro Sanchez Gonzalez, Peter Battaglia, Jessica B Hamrick, and Tobias Pfaff. Inverse design for fluid-structure interactions using graph network simulators. *Advances in Neural Information Processing Systems*, 35:13759–13774, 2022.
- [138] Luis Tenorio. *An introduction to data analysis and uncertainty quantification for inverse problems*. SIAM, 2017.
- [139] Joseph B Nagel and Bruno Sudret. A unified framework for multilevel uncertainty quantification in bayesian inverse problems. *Probabilistic Engineering Mechanics*, 43: 68–84, 2016.
- [140] Marco Iglesias and Andrew M Stuart. Inverse problems and uncertainty quantification. *SIAM News*, 20:2–3, 2014.
- [141] Bruno Sudret, Stefano Marelli, and Joe Wiart. Surrogate models for uncertainty quantification: An overview. In *2017 11th European conference on antennas and propagation (EUCAP)*, pages 793–797. IEEE, 2017.

- [142] Alexander Forrester, Andras Sobester, and Andy Keane. *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons, 2008.
- [143] Bruno Sudret. Surrogate models for uncertainty quantification in engineering sciences. In *Dutch Computational Science Day (DUCOMS 2024)*. ETH Zurich, Chair of Risk, Safety and Uncertainty Quantification, 2024.
- [144] David JC MacKay et al. Introduction to gaussian processes. *NATO ASI series F computer and systems sciences*, 168:133–166, 1998.
- [145] Christopher Williams and Carl Rasmussen. Gaussian processes for regression. *Advances in neural information processing systems*, 8, 1995.
- [146] Eric Schulz, Maarten Speekenbrink, and Andreas Krause. A tutorial on gaussian process regression: Modelling, exploring, and exploiting functions. *Journal of mathematical psychology*, 85:1–16, 2018.
- [147] Robert B Gramacy. *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. Chapman and Hall/CRC, 2020.
- [148] Jinglai Li and Youssef M Marzouk. Adaptive construction of surrogates for the bayesian solution of inverse problems. *SIAM Journal on Scientific Computing*, 36(3):A1163–A1186, 2014.
- [149] Michalis Frangos, Youssef Marzouk, Karen Willcox, and Bart van Bloemen Waanders. Surrogate and reduced-order modeling: a comparison of approaches for large-scale statistical inverse problems. *Large-Scale Inverse Problems and Quantification of Uncertainty*, pages 123–149, 2010.
- [150] Bowen Lei, Tanner Quinn Kirk, Anirban Bhattacharya, Debdeep Pati, Xiaoning Qian, Raymundo Arroyave, and Bani K Mallick. Bayesian optimization with adaptive surrogate models for automated experimental design. *Npj Computational Materials*, 7(1):194, 2021.
- [151] Bruno Sudret. Surrogate models for forward and inverse uncertainty quantification. In *SSD Seminar Series*. ETH Zurich, Chair of Risk, Safety and Uncertainty Quantification, 2021.
- [152] Dirk Gorissen, Ivo Couckuyt, Piet Demeester, Tom Dhaene, and Karel Crombecq. A surrogate modeling and adaptive sampling toolbox for computer based design. *The Journal of Machine Learning Research*, 11:2051–2055, 2010.
- [153] Lisia Dias, Atharv Bhosekar, and Mariathi Ierapetritou. Adaptive sampling approaches for surrogate-based optimization. In *Computer aided chemical engineering*, volume 47, pages 377–384. Elsevier, 2019.

- [154] Aliakbar Golzari, Morteza Haghghat Sefat, and Saeid Jamshidi. Development of an adaptive surrogate model for production optimization. *Journal of petroleum Science and Engineering*, 133:677–688, 2015.
- [155] Bruno Gaspar, Angelo P Teixeira, and C Guedes Soares. Adaptive surrogate model with active refinement combining kriging and a trust region method. *Reliability Engineering & System Safety*, 165:277–291, 2017.
- [156] Ben G Fitzpatrick. Bayesian analysis in inverse problems. *Inverse problems*, 7(5):675, 1991.
- [157] Jérôme Idier. *Bayesian approach to inverse problems*. John Wiley & Sons, 2013.
- [158] Christophe Andrieu and Johannes Thoms. A tutorial on adaptive mcmc. *Statistics and computing*, 18(4):343–373, 2008.
- [159] Christian P Robert and George Casella. Monte carlo integration. In *Monte Carlo statistical methods*, pages 71–138. Springer, 1999.
- [160] Dave Higdon, James Gattiker, Brian Williams, and Maria Rightley. Computer model calibration using high-dimensional output. *Journal of the American Statistical Association*, 103(482):570–583, 2008.
- [161] Youssef M Marzouk, Habib N Najm, and Larry A Rahn. Stochastic spectral methods for efficient bayesian solution of inverse problems. *Journal of Computational Physics*, 224(2): 560–586, 2007.
- [162] Bartek T Knapik, Aad W Van Der Vaart, and J Harry van Zanten. Bayesian inverse problems with gaussian priors. 2011.
- [163] Heber L Rocha, João Vitor de O Silva, Renato S Silva, Ernesto ABF Lima, and Regina C Almeida. Bayesian inference using gaussian process surrogates in cancer modeling. *Computer Methods in Applied Mechanics and Engineering*, 399:115412, 2022.
- [164] Jonas El Gammal, Nils Schöneberg, Jesús Torrado, and Christian Fidler. Fast and robust bayesian inference using gaussian processes with gprry. *Journal of Cosmology and Astroparticle Physics*, 2023(10):021, 2023.
- [165] Vladimir Spokoiny. Inexact approximation and the use of posterior mean in bayesian inference. *Bayesian Analysis*, 20(1):1–28, 2025.
- [166] Wilhelm Werner. Polynomial interpolation: Lagrange versus newton. *Mathematics of computation*, pages 205–217, 1984.
- [167] Bradley K Alpert and Vladimir Rokhlin. A fast algorithm for the evaluation of legendre expansions. *SIAM Journal on Scientific and Statistical Computing*, 12(1):158–179, 1991.

- [168] Sky McKinley and Megan Levine. Cubic spline interpolation. *College of the Redwoods*, 45(1):1049–1060, 1998.
- [169] Andrew Gordon Wilson, David A Knowles, and Zoubin Ghahramani. Gaussian process regression networks. *arXiv preprint arXiv:1110.4411*, 2011.
- [170] Arman Melkumyan and Fabio Ramos. Multi-kernel gaussian processes. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, page 1408, 2011.
- [171] B Gaspar, Angelo Palos Teixeira, and C Guedes Soares. A study on a stopping criterion for active refinement algorithms in kriging surrogate models. *Safety and reliability of complex engineered systems*, pages 1219–1227, 2015.
- [172] Mihaela Chiappetta. UQforPython: Uncertainty quantification for python, 2026. URL <https://github.com/mihaelachiappetta/UQforPython>. GitHub repository, accessed 2026-03-02.
- [173] Simon L Cotter, Massoumeh Dashti, James Cooper Robinson, and Andrew M Stuart. Bayesian inverse problems for functions and applications to fluid mechanics. *Inverse problems*, 25(11):115008, 2009.
- [174] Gary D Knott. *Interpolating cubic splines*, volume 18. Springer Science & Business Media, 1999.
- [175] Habib N Najm. Uncertainty quantification and polynomial chaos techniques in computational fluid dynamics. *Annual review of fluid mechanics*, 41(1):35–52, 2009.
- [176] Martin Dietrich Buhmann. Radial basis functions. *Acta numerica*, 9:1–38, 2000.
- [177] Bengt Fornberg and Julia Zuev. The Runge phenomenon and spatially variable shape parameters in RBF interpolation. *Computers & Mathematics with Applications*, 54(3):379–398, 2007.
- [178] Charles W Groetsch and CW Groetsch. *Inverse problems in the mathematical sciences*, volume 52. Springer, 1993.
- [179] Nikolaos Kladovasilakis, Konstantinos Tsongas, Dimitris Karalekas, and Dimitrios Tzetzis. Architected materials for additive manufacturing: a comprehensive review. *Materials*, 15(17):5919, 2022.
- [180] Mikhail Osanov and James K Guest. Topology optimization for architected materials design. *Annual Review of Materials Research*, 46(1):211–233, 2016.
- [181] Marco Maurizi, Chao Gao, and Filippo Berto. Inverse design of truss lattice materials with superior buckling resistance. *npj Computational Materials*, 8(1):247, 2022.

- [182] Chan Soo Ha, Desheng Yao, Zhenpeng Xu, Chenang Liu, Han Liu, Daniel Elkins, Matthew Kile, Vikram Deshpande, Zhenyu Kong, Mathieu Bauchy, et al. Rapid inverse design of metamaterials based on prescribed mechanical behavior through machine learning. *Nature Communications*, 14(1):5765, 2023.
- [183] Grégoire Allaire and Robert Brizzi. A multiscale finite element method for numerical homogenization. *Multiscale Modeling & Simulation*, 4(3):790–812, 2005.
- [184] VG Kouznetsova, MG Geers, and WAM Brekelmans. An approach to micro-macro modeling of heterogeneous materials. *Computational Mechanics*, 27:37–48, 2001.
- [185] Nathan O Siu and Dana L Kelly. Bayesian parameter estimation in probabilistic risk assessment. *Reliability Engineering & System Safety*, 62(1-2):89–116, 1998.
- [186] Carla Currin, Toby Mitchell, Max Morris, and Don Ylvisaker. Bayesian prediction of deterministic functions, with applications to the design and analysis of computer experiments. *Journal of the American Statistical Association*, 86(416):953–963, 1991.
- [187] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [188] Amandine Marrel and Bertrand Iooss. Probabilistic surrogate modeling by gaussian process: A review on recent insights in estimation and validation. *Reliability Engineering & System Safety*, 247:110094, 2024.
- [189] S Marelli, C Lamas, K Konakli, C Mylonas, P Wiederkehr, and B Sudret. Uqlab user manual—sensitivity analysis. *Report UQLab-VI*, pages 2–106, 2019.
- [190] David D Jackson and Mitsuhiro Matsu’Ura. A bayesian approach to nonlinear inversion. *Journal of Geophysical Research: Solid Earth*, 90(B1):581–591, 1985.
- [191] Hermann G Matthies, Elmar Zander, Bojana V Rosić, Alexander Litvinenko, and Oliver Pajonk. Inverse problems in a bayesian setting. In *Computational Methods for Solids and Fluids: Multiscale Analysis, Probability Aspects and Model Reduction*, pages 245–286. Springer, 2016.
- [192] Dave Higdon, Marc Kennedy, James C Cavendish, John A Cafeo, and Robert D Ryne. Combining field data and computer simulations for calibration and prediction. *SIAM Journal on Scientific Computing*, 26(2):448–466, 2004.
- [193] Liang Yan and Yuan-Xiang Zhang. Convergence analysis of surrogate-based methods for bayesian inverse problems. *Inverse Problems*, 33(12):125001, 2017.
- [194] Haitao Liu, Yew-Soon Ong, Xiaobo Shen, and Jianfei Cai. When gaussian process meets big data: A review of scalable gps. *IEEE transactions on neural networks and learning systems*, 31(11):4405–4423, 2020.

- [195] Sergey Oladyshkin, Farid Mohammadi, Ilja Kroeker, and Wolfgang Nowak. Bayesian active learning for the gaussian process emulator using information theory. *Entropy*, 22(8):890, 2020.
- [196] Michael J Asher, Barry FW Croke, Anthony J Jakeman, and Luk JM Peeters. A review of surrogate models and their application to groundwater modeling. *Water Resources Research*, 51(8):5957–5973, 2015.
- [197] Vincent Dubourg. *Adaptive surrogate models for reliability analysis and reliability-based design optimization*. PhD thesis, Université Blaise Pascal-Clermont-Ferrand II, 2011.
- [198] Alexander Raßloff, Paul Seibert, Karl A Kalina, and Markus Kästner. Inverse design of spinodoid structures using bayesian optimization. *Computational Mechanics*, 77(1):275–296, 2026.
- [199] Max Rosenkranz, Markus Kästner, and Ivo F Sbalzarini. Data-efficient inverse design of spinodoid metamaterials. *Integrating Materials and Manufacturing Innovation*, 14(4):695–713, 2025.
- [200] Julien Yvonnet. *Computational homogenization of heterogeneous materials with finite elements*, volume 258. Springer, 2019.
- [201] John W Cahn. On spinodal decomposition. *Acta metallurgica*, 9(9):795–801, 1961.
- [202] V Sofonea and KR Mecke. Morphological characterization of spinodal decomposition kinetics. *The European Physical Journal B-Condensed Matter and Complex Systems*, 8(1):99–112, 1999.
- [203] Christopher P Grant. Spinodal decomposition for the cahn-hilliard equation. *Communications in Partial Differential Equations*, 18(3-4):453–490, 1993.
- [204] Hidde Derk Vuijk, Joseph Michael Brader, and Abhinav Sharma. Effect of anisotropic diffusion on spinodal decomposition. *Soft matter*, 15(6):1319–1326, 2019.
- [205] Ananthan Vidyasagar, Sebastian Krödel, and Dennis M Kochmann. Microstructural patterns with tunable mechanical anisotropy obtained by simulating anisotropic spinodal decomposition. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2218), 2018.
- [206] Meng-Ting Hsieh, Bianca Endo, Yunfei Zhang, Jens Bauer, and Lorenzo Valdevit. The mechanical response of cellular materials with spinodal topologies. *Journal of the Mechanics and Physics of Solids*, 125:401–419, 2019.
- [207] Robert J Adler and Jonathan E Taylor. *Random fields and geometry*. Springer, 2007.

- [208] Data Analytics Engineering Group. Fans: Fourier-accelerated nodal solver. <https://github.com/DataAnalyticsEngineering/FANS/>, 2025. Accessed: 2026-04-02.
- [209] Matthias Leuschner and Felix Fritzen. Fourier-accelerated nodal solvers (fans) for homogenization problems. *Computational Mechanics*, 62(3):359–392, 2018.
- [210] Sanath Keshav, Felix Fritzen, and Matthias Kabel. Fft-based homogenization at finite strains using composite boxels (combo). *Computational Mechanics*, 71(1):191–212, 2023.
- [211] Martin Ladecký, Richard J Leute, Ali Falsafi, Ivana Pultarová, Lars Pastewka, Till Junge, and Jan Zeman. Optimal fft-accelerated finite element solver for homogenization. *arXiv preprint arXiv:2203.02962*, 2022.
- [212] Alphonse Finel. A fast and robust discrete fft-based solver for computational homogenization. *arXiv preprint arXiv:2405.11168*, 2024.
- [213] Sébastien Brisard and Luc Dormieux. Fft-based methods for the mechanics of composites: A general variational framework. *Computational Materials Science*, 49(3):663–671, 2010.
- [214] Joseph B Nagel. Bayesian techniques for inverse uncertainty quantification. *IBK Bericht*, 504, 2019.
- [215] Maria Jesús Bayarri, James O Berger, and Fei Liu. Modularization in bayesian analysis, with emphasis on analysis of computer models. 2009.
- [216] Jie Zhang, Souma Chowdhury, Ali Mehmani, and Achille Messac. Characterizing uncertainty attributable to surrogate models. *Journal of Mechanical Design*, 136(3):031004, 2014.
- [217] Carl Edward Rasmussen. Advances in gaussian processes. *Advances in Neural Information Processing Systems*, 19, 2006.
- [218] Gerhard Venter. Review of optimization techniques. 2010.
- [219] Hans A Eschenauer and Niels Olhoff. Topology optimization of continuum structures: a review. *Appl. Mech. Rev.*, 54(4):331–390, 2001.
- [220] Martin P Bendsøe. *Optimization of structural topology, shape, and material*, volume 414. Springer, 1995.
- [221] Michael Bremicker, Mehran Chirehdast, Noboru Kikuchi, and Panos Y Papalambros. Integrated topology and shape optimization in structural design. *Journal of Structural Mechanics*, 19(4):551–587, 1991.
- [222] Loris Barbieri and Maurizio Muzzupappa. Performance-driven engineering design approaches based on generative design and topology optimization tools: a comparative study. *Applied Sciences*, 12(4):2106, 2022.

- [223] Martin Philip Bendsøe and Noboru Kikuchi. Generating optimal topologies in structural design using a homogenization method. *Computer methods in applied mechanics and engineering*, 71(2):197–224, 1988.
- [224] Martin P Bendsøe. Optimal shape design as a material distribution problem. *Structural optimization*, 1:193–202, 1989.
- [225] Nico P Van Dijk, Kurt Maute, Matthijs Langelaar, and Fred Van Keulen. Level-set methods for structural topology optimization: a review. *Structural and Multidisciplinary Optimization*, 48:437–472, 2013.
- [226] Julian A Norato, Martin P Bendsøe, Robert B Haber, and Daniel A Tortorelli. A topological derivative method for topology optimization. *Structural and Multidisciplinary Optimization*, 33:375–386, 2007.
- [227] Massimo Carraturo, Elisabetta Rocca, Elena Bonetti, Dietmar Hömberg, Alessandro Reali, and Ferdinando Auricchio. Graded-material design based on phase-field and topology optimization. *Computational Mechanics*, 64:1589–1600, 2019.
- [228] Xiaodong Huang and Mike Xie. *Evolutionary topology optimization of continuum structures: methods and applications*. John Wiley & Sons, 2010.
- [229] Mariusz Bujny, Nikola Aulig, Markus Olhofer, and Fabian Duddeck. Hybrid evolutionary approach for level set topology optimization. In *2016 IEEE Congress on Evolutionary Computation (CEC)*, pages 5092–5099. IEEE, 2016.
- [230] Gouri Dhatt, Emmanuel Lefrançois, and Gilbert Touzot. *Finite element method*. John Wiley & Sons, 2012.
- [231] Behrooz Hassani and Ernest Hinton. A review of homogenization and topology optimization iii—topology optimization using optimality criteria. *Computers & structures*, 69(6): 739–756, 1998.
- [232] GIN Rozvany and M Zhou. Continuum-based optimality criteria (coc) methods: an introduction. In *Optimization of large structural systems*, pages 1–26. Springer, 1993.
- [233] Antony Jameson. Gradient based optimization methods. *MAE Technical Report No*, (2057), 1995.
- [234] Singiresu S Rao. *Engineering optimization: theory and practice*. John Wiley & Sons, 2019.
- [235] Gilles Marck, Maroun Nemer, Jean-Luc Harion, Serge Russeil, and Daniel Bougeard. Topology optimization using the SIMP method for multiobjective conductive problems. *Numerical Heat Transfer, Part B: Fundamentals*, 61(6):439–470, 2012.

- [236] Dongruo Zhou, Jinghui Chen, Yuan Cao, Ziyang Yang, and Quanquan Gu. On the convergence of adaptive gradient methods for nonconvex optimization. *arXiv preprint arXiv:1808.05671*, 2018.
- [237] Krister Svanberg. The Method of Moving Asymptotes-Modelling aspects and solution schemes. *Lecture Notes for the DCAMM course Advanced Topics in Structural Optimization*, 1998.
- [238] Krister Svanberg. MMA and GCMMA-two methods for nonlinear optimization. *vol*, 1: 1–15, 2007.
- [239] Krister Svanberg. MMA and GCMMA–Fortran versions March 2013. *Royal Institute of Technology*, pages 1–23, 2013.
- [240] Niels Aage, Erik Andreassen, and Boyan Stefanov Lazarov. Topology optimization using PETS: An easy-to-use, fully parallel, open source topology optimization framework. *Structural and Multidisciplinary Optimization*, 51(3):565–572, 2015.
- [241] Shi-An Zhou and Song Yao. Stress-based topology optimization of continuum structures incorporating a piecewise P-norm stabilization strategy. *Finite Elements in Analysis and Design*, 250:104401, 2025.
- [242] Michael Bruyneel, Pierre Duysinx, and Claude Fleury. A family of MMA approximations for structural optimization. *Structural and Multidisciplinary Optimization*, 24(4):263–276, 2002.
- [243] Kong-Tian Zuo, Li-Ping Chen, Yun-Qing Zhang, and Jingzhou Yang. A hybrid topology optimization algorithm for structural design. *Engineering Optimization*, 37(8):849–866, 2005.
- [244] Marco Montemurro and Khalil Refai. A topology optimization method based on non-uniform rational basis spline hyper-surfaces for heat conduction problems. *Symmetry*, 13(5):888, 2021.
- [245] Blaise Bourdin. Filters in topology optimization. *International journal for numerical methods in engineering*, 50(9):2143–2158, 2001.
- [246] Shengyin Wang and Michael Yu Wang. Radial basis functions and level set method for structural topology optimization. *International journal for numerical methods in engineering*, 65(12):2060–2090, 2006.
- [247] Robert B Haber, Chandrashekhar S Jog, and Martin P Bendsøe. A new approach to variable-topology shape design using a constraint on perimeter. *Structural optimization*, 11:1–12, 1996.

- [248] Thomas Borrvall and Joakim Petersson. Topology optimization using regularized intermediate density control. *Computer Methods in Applied Mechanics and Engineering*, 190 (37-38):4911–4928, 2001.
- [249] M Zhou, YK Shyy, and HL Thomas. Checkerboard and minimum member size control in topology optimization. *Structural and Multidisciplinary Optimization*, 21:152–158, 2001.
- [250] Avinash Shukla, Anadi Misra, and Sunil Kumar. Checkerboard problem in finite element based topology optimization. *International Journal of Advances in Engineering & Technology*, 6(4):1769, 2013.
- [251] Kai Liu and Andrés Tovar. An efficient 3d topology optimization code written in matlab. *Structural and multidisciplinary optimization*, 50:1175–1196, 2014.
- [252] Godfrey C Onwubolu and BV Babu. *New optimization techniques in engineering*, volume 141. Springer, 2013.
- [253] R Venkata Rao and Vimal J Savsani. Advanced optimization techniques. In *Mechanical Design Optimization Using Advanced Optimization Techniques*, pages 5–34. Springer, 2012.
- [254] Hubertus Th Jongen, Klaus Meer, and Eberhard Triesch. *Optimization theory*. Springer, 2004.
- [255] Kalyanmoy Deb. *Optimization for engineering design: Algorithms and examples*. PHI Learning Pvt. Ltd., 2012.
- [256] Leslie R Foulds. *Optimization techniques: an introduction*. Springer Science & Business Media, 2012.
- [257] Andrzej Ruszczyński. *Nonlinear optimization*. Princeton university press, 2011.
- [258] Didier Henrion and J-B Lasserre. Solving nonconvex optimization problems. *IEEE Control Systems Magazine*, 24(3):72–83, 2004.
- [259] Xiaohui Hu, Russell Eberhart, et al. Solving constrained nonlinear optimization problems with particle swarm optimization. In *Proceedings of the sixth world multiconference on systemics, cybernetics and informatics*, volume 5, pages 203–206. Citeseer, 2002.
- [260] Boris Sholimovich Mordukhovich. *Variational analysis and applications*. Springer, 2018.
- [261] Peter W Christensen and Anders Klarbring. Sequential explicit, convex approximations. In *An Introduction to Structural Optimization*, pages 57–75. Springer, 2008.
- [262] Claude Fleury. Sequential convex programming for structural optimization problems. In *Optimization of large structural systems*, pages 531–553. Springer, 1993.

- [263] Amir Beck. *First-order methods in optimization*. SIAM, 2017.
- [264] Mohammad Sh Daoud, Mohammad Shehab, Hani M Al-Mimi, Laith Abualigah, Raed Abu Zitar, and Mohd Khaled Yousef Shambour. Gradient-based optimizer (GBO): a review, theory, variants, and applications. *Archives of Computational Methods in Engineering*, 30(4):2431–2449, 2023.
- [265] Saleh Masoud Abdallah Altbawi, Saifulnizam Bin Abdul Khalid, Ahmad Safawi Bin Mokhtar, Hussain Shareef, Nusrat Husain, Ashraf Yahya, Syed Aqeel Haider, Lubna Moin, and Rayan Hamza Alsisi. An improved gradient-based optimization algorithm for solving complex optimization problems. *Processes*, 11(2):498, 2023.
- [266] Manzil Zaheer, Sashank Reddi, Devendra Sachan, Satyen Kale, and Sanjiv Kumar. Adaptive methods for nonconvex optimization. *Advances in neural information processing systems*, 31, 2018.
- [267] Xiaoxia Wu, Simon S Du, and Rachel Ward. Global convergence of adaptive gradient methods for an over-parameterized neural network. *arXiv preprint arXiv:1902.07111*, 2019.
- [268] Grzegorz W. Wasilkowski and Henryk Wozniakowski. Explicit Cost Bounds of Algorithms for Multivariate Tensor Product Problems. *Journal of Complexity*, 11(1):1–56, 1995.

Appendix A

Sparse-grid surrogate modeling

In the present section, the sparse-grid surrogate modeling approach adopted for the study is presented. Following the notation introduced in Section 2.5, the problem of approximating an N -variate scalar function $f(\mathbf{v}) : \Gamma \rightarrow \mathbb{R}$ is considered, where $\mathbf{v} \in \Gamma \subset \mathbb{R}^N$ (extension to P -valued functions $\mathbf{f} : \Gamma \rightarrow \mathbb{R}^P$ is immediate; it is enough to apply the same procedure to each component of \mathbf{f}). It is also recalled that v_n are independent random variables with probability density function $\rho_n(v_n)$, $n = 1, \dots, N$ and that therefore the joint probability density of \mathbf{v} over Γ is the product $\rho(\mathbf{v}) = \prod_{n=1}^N \rho_n(v_n)$.

The first step in constructing the sparse-grid surrogate model is to define a set of collocation points for each parameter v_n . The number of points along v_n is denoted by $K_n \in \mathbb{N}_+$, and a discretization level for each parameter is defined, i.e., a positive number $i_n \in \mathbb{N}_+$, $i_n \geq 1$, using a “level-to-knots” function m that associates to each level a number of points:

$$m : \mathbb{N}_+ \rightarrow \mathbb{N}_+ \text{ such that } m(i_n) = K_n. \quad (\text{A.1})$$

In this work, $m(i_n) = 2i_n - 1$ is considered (i.e., at each level i_n two more points with respect to the previous level are considered; cf. Equation (2.14)), but other choices are possible. The set of collocation points at level i_n along parameter v_n is denoted by:

$$\mathcal{T}_{i_n} = \left\{ y_{n,m(i_n)}^{(j_n)} : j_n = 1, \dots, m(i_n) \right\} \quad \text{for } n = 1, \dots, N. \quad (\text{A.2})$$

The positions of these points over Γ_n are usually chosen on the basis of the PDF ρ_n of the random variables v_n . As reported in Table 2.7, in this work symmetric Leja points are used whenever v_n is a uniform random variable and symmetric Gaussian Leja points whenever v_n is a Gaussian random variable (see B for details), but other choices are possible, see e.g. [80]. Such choices have the advantage that Leja points are nested, i.e., $\mathcal{T}_{i_n} \subset \mathcal{T}_{l_n}$ if $l_n > i_n$.

The second step is the definition of tensor grids of N dimensions, derived as the Cartesian product of the previously introduced univariate sets \mathcal{T}_{i_n} , and of their associated Lagrangian interpolants. In particular, by collecting the discretization levels i_n in a multi-index $\mathbf{i} \in \mathbb{N}_+^N$, considering the corresponding tensor grid $\overline{\mathcal{T}}_{\mathbf{i}} = \bigotimes_{n=1}^N \mathcal{T}_{i_n}$, with number of nodes $M_{\mathbf{i}} = \prod_{n=1}^N m(i_n)$ it is possible to write:

$$\overline{\mathcal{T}}_{\mathbf{i}} = \left\{ \mathbf{v}_{m(\mathbf{i})}^{(\mathbf{j})} \right\}_{\mathbf{j} \leq m(\mathbf{i})}, \quad \text{with} \quad \mathbf{v}_{m(\mathbf{i})}^{(\mathbf{j})} = \left[v_{1,m(i_1)}^{(j_1)}, \dots, v_{N,m(i_N)}^{(j_N)} \right] \text{ and } \mathbf{j} \in \mathbb{N}_+^N,$$

where $m(\mathbf{i}) = [m(i_1), m(i_2), \dots, m(i_N)]$ and $\mathbf{j} \leq m(\mathbf{i})$ means that $j_n \leq m(i_n)$ for every $n = 1, \dots, N$. The tensor-interpolant approximation (also called tensor-interpolant *surrogate*

model) of $f(\mathbf{v})$, denoted by $\mathcal{U}_{\mathbf{i}}(\mathbf{v})$, is then an N -variate Lagrangian interpolant collocated at the grid nodes of $\mathcal{T}_{\mathbf{i}}$ and can be written as:

$$f(\mathbf{v}) \approx \mathcal{U}_{\mathbf{i}}(\mathbf{v}) := \sum_{\mathbf{j} \leq \mathbf{i}} f\left(v_{m(\mathbf{i})}^{(\mathbf{j})}\right) \mathcal{L}_{m(\mathbf{i})}^{(\mathbf{j})}(\mathbf{v}), \quad (\text{A.3})$$

where $\left\{ \mathcal{L}_{m(\mathbf{i})}^{(\mathbf{j})}(\mathbf{v}) \right\}_{\mathbf{j} \leq \mathbf{i}}$ are N -variate Lagrange polynomials, defined as tensor products of univariate Lagrange polynomials, i.e.

$$\mathcal{L}_{m(\mathbf{i})}^{(\mathbf{j})}(\mathbf{v}) = \prod_{n=1}^N \ell_{n,m(i_n)}^{(j_n)}(v_n) \quad \text{with} \quad \ell_{n,m(i_n)}^{(j_n)}(v_n) = \prod_{k=1, k \neq j_n}^{m(i_n)} \frac{v_n - v_{n,m(i_n)}^{(k)}}{v_{n,m(i_n)}^{(k)} - v_{n,m(i_n)}^{(j_n)}}.$$

The accuracy of the approximation $f(\mathbf{v}) \approx \mathcal{U}_{\mathbf{i}}(\mathbf{v})$ increases as the number of collocation points in each v_n grows, i.e., for $i_n \gg 1$, $n = 1, \dots, N$. At the same time, the cost of constructing $\mathcal{U}_{\mathbf{i}}(\mathbf{v})$ grows exponentially in N , since it requires evaluating f at $M_{\mathbf{i}} = \prod_{n=1}^N m(i_n)$ points; this implies that even moderate choices of i_n , $n = 1, \dots, N$ could be unfeasible for $N > 2$ if evaluating f is an expensive operation. To mitigate this problem, the sparse-grid surrogate model consists of an approximation of $f(\mathbf{v})$ formed by a linear combination of several coarse $\mathcal{U}_{\mathbf{i}}(\mathbf{v})$ rather than by a single $\mathcal{U}_{\mathbf{i}}(\mathbf{v})$ with $i_n \gg 1$, $n = 1, \dots, N$.

For this purpose, as a third step towards sparse-grid surrogate models the so-called univariate and multivariate detail operators are introduced:

$$\Delta_n[\mathcal{U}_{\mathbf{i}}(\mathbf{v})] = \mathcal{U}_{\mathbf{i}}(\mathbf{v}) - \mathcal{U}_{\mathbf{i}-\mathbf{e}_n}(\mathbf{v}) \quad \text{with} \quad 1 \leq n \leq N; \quad (\text{A.4})$$

$$\Delta[\mathcal{U}_{\mathbf{i}}(\mathbf{v})] = \bigotimes_{n=1}^N \Delta_n[\mathcal{U}_{\mathbf{i}}(\mathbf{v})] = \Delta_1[\dots[\Delta_N[\mathcal{U}_{\mathbf{i}}(\mathbf{v})]]], \quad (\text{A.5})$$

where $\mathcal{U}_{\mathbf{i}}(\mathbf{v}) = 0$ when at least one component of \mathbf{i} is zero and \mathbf{e}_n the n -th canonical multi-index, i.e., $(\mathbf{e}_n)_k = 1$ if $n = k$ and 0 otherwise. Multivariate detail operators can be evaluated as suitable linear combinations of certain approximations of the complete tensor approximations $\mathcal{U}_{\mathbf{i}}$:

$$\Delta[\mathcal{U}_{\mathbf{i}}(\mathbf{v})] = \Delta_1[\dots[\Delta_N[\mathcal{U}_{\mathbf{i}}]]] = \sum_{\mathbf{j} \in \{0,1\}^N} (-1)^{\|\mathbf{j}\|_1} \mathcal{U}_{\mathbf{i}-\mathbf{j}}(\mathbf{v}). \quad (\text{A.6})$$

Moreover, note that a hierarchical decomposition of $\mathcal{U}_{\mathbf{i}}(\mathbf{v})$ holds:

$$\mathcal{U}_{\mathbf{i}}(\mathbf{v}) = \sum_{\mathbf{j} \leq \mathbf{i}} \Delta[\mathcal{U}_{\mathbf{j}}(\mathbf{v})]. \quad (\text{A.7})$$

The fourth and final step to construct a sparse-grid surrogate model is to tweak such hierarchical decomposition. In detail, instead of summing over $\mathbf{j} \leq \mathbf{i}$ a different collection of multi-indices \mathcal{I} (from here on, multi-index set) is considered, chosen according to criteria that will be made clearer in a moment:

$$f(\mathbf{v}) \approx \mathcal{S}_{\mathcal{I}} f(\mathbf{v}) = \sum_{\mathbf{i} \in \mathcal{I}} \Delta[\mathcal{U}_{\mathbf{i}}(\mathbf{v})].$$

Furthermore, applying Equation (A.6) a more practical expression is obtained, i.e., the so-called “combination technique” [268], which is the form actually implemented in the Sparse-Grids Matlab-Kit:

$$f(\mathbf{v}) \approx \mathcal{S}_{\mathcal{I}} f(\mathbf{v}) = \sum_{\mathbf{i} \in \mathcal{I}} c_{\mathbf{i}} \mathcal{U}_{\mathbf{i}}(\mathbf{v}), \quad c_{\mathbf{i}} := \sum_{\substack{\mathbf{j} \in \{0,1\}^N \\ \mathbf{i} + \mathbf{j} \in \mathcal{I}}} (-1)^{\|\mathbf{j}\|_1}. \quad (\text{A.8})$$

This re-writing is valid only if \mathcal{I} is *downward-closed*, i.e., if \mathcal{I} is such that if a certain multi-index \mathbf{i} is in \mathcal{I} all its “previous” multi-indices $\mathbf{j} \leq \mathbf{i}$ are also in the set. In formulae, this requires:

$$\forall \mathbf{i} \in \mathcal{I}, \mathbf{i} - \mathbf{e}_n \in \mathcal{I}, \quad \forall n = 1, \dots, N \text{ s.t. } i_n > 0. \quad (\text{A.9})$$

Concerning the choice of the multi-index set \mathcal{I} , the idea is to discard from the hierarchical decomposition in Equation (A.7) the contributions that have a large cost and contribute little to the approximation (in a sense, dropping the high-order corrections). Under mild regularity assumptions of $f(\mathbf{v})$, a simple yet effective choice to this end is

$$\mathcal{I}_{sum} = \{\mathbf{i} \in \mathbb{N}_+^N : \sum_{n=1}^N (i_n - 1) \leq w\} \quad (\text{A.10})$$

for some integer value w (the larger w , the more accurate is the sparse-grid surrogate model). Note that conversely, choosing

$$\mathcal{I}_{max} = \{\mathbf{i} \in \mathbb{N}_+^N : \max_{n=1, \dots, N} (i_n - 1) \leq w\} \quad (\text{A.11})$$

a tensor grid with $m(w + 1)$ points per direction is obtained, i.e., $\mathcal{S}_{\mathcal{I}_{max}(w)}(\mathbf{v}) = \mathcal{U}_{\mathbf{i}_w}(\mathbf{v})$ with $\mathbf{i}_w = [w+1, w+1, \dots]$; this is an immediate consequence of the decomposition in Equation (A.7). More advanced options to tailor the set \mathcal{I} to the function f are available in literature, and in particular it would be possible to use an adaptive algorithm, that adds multi-indices \mathbf{i} to \mathcal{I} one by one given the values of f ; see again [80] for details. Finally, the term *sparse grid* denotes the collection of points needed to build the *sparse-grid surrogate model* $\mathcal{S}_{\mathcal{I}}$, i.e.

$$\mathcal{G}_{\mathcal{I}} = \bigcup_{\substack{\mathbf{i} \in \mathcal{I} \\ c_{\mathbf{i}} \neq 0}} \mathcal{T}_{\mathbf{i}}. \quad (\text{A.12})$$

Appendix B

Leja points

Leja knots have been introduced for unweighted interpolation on intervals $[a, b]$, see [80, 123, 124] and references therein, and are therefore a suitable choice when v_n are uniform random variables. They are built recursively as:

$$v_n^{(1)} = b, \quad v_n^{(2)} = a, \quad v_n^{(3)} = \frac{a+b}{2}, \quad v_n^{(j)} = \arg \max_{v_n \in [a, b]} \prod_{k=1}^{j-1} |v_n - v_n^{(k)}|. \quad (\text{B.1})$$

By construction Leja knots are nested but not symmetric with respect to the mid-point $\frac{a+b}{2}$, which is also a desirable property. To fix this issue, the construction above can be changed by generating only the even elements of the sequence with the standard formula in (Equation (B.1)) and then symmetrizing them to obtain the odd elements, i.e.

$$\begin{aligned} v_n^{(1)} &= b, \quad v_n^{(2)} = a, \quad v_n^{(3)} = \frac{a+b}{2}, \\ v_n^{(2j)} &= \arg \max_{v_n \in [a, b]} \prod_{k=1}^{2j-1} |v_n - v_n^{(k)}|, \\ v_n^{(2j+1)} &= \frac{a+b}{2} - \left(v_n^{(2j)} - \frac{a+b}{2} \right). \end{aligned} \quad (\text{B.2})$$

It is furthermore possible to extend the construction of Leja knots to the case when $v_n \in \Gamma_n$ are Gaussian random variables (or more generally, random variables with a probability distribution other than uniform), see again [80, 123, 124]. The knots thus obtained are the so-called Gaussian Leja knots (or in general, weighted Leja knots) and can be computed again recursively, by suitably introducing a weight in Equation (B.1), i.e., solving

$$v_n^{(j)} = \arg \max_{v_n \in \Gamma_n} \sqrt{\rho_n(v_n)} \prod_{k=1}^{j-1} |v_n - v_n^{(k)}|,$$

where ρ_n is the PDF of the random variable. Symmetric versions of Gaussian (weighted) Leja points can then be generated following the procedure that leads to Equation (B.2).

Appendix C

Implementation Details: The UQforPy Repository

All numerical experiments presented in this work are implemented within the open-source Python repository UQforPy [172]. The repository provides a modular computational framework for uncertainty quantification, surrogate modeling, and inverse problem analysis, and has been specifically designed to support reproducible research in the context of surrogate-assisted Bayesian methodologies for computationally expensive models.

The software follows a structured scientific Python architecture and is organized around the directory `src/uqforpy`, which contains the core implementation of the framework. The overall organization of the repository is illustrated in Figure C.1. The implemented components reflect the main methodological building blocks considered in the present work, including sampling strategies, surrogate modeling, Bayesian optimization, Bayesian inversion, uncertainty quantification, and global sensitivity analysis.

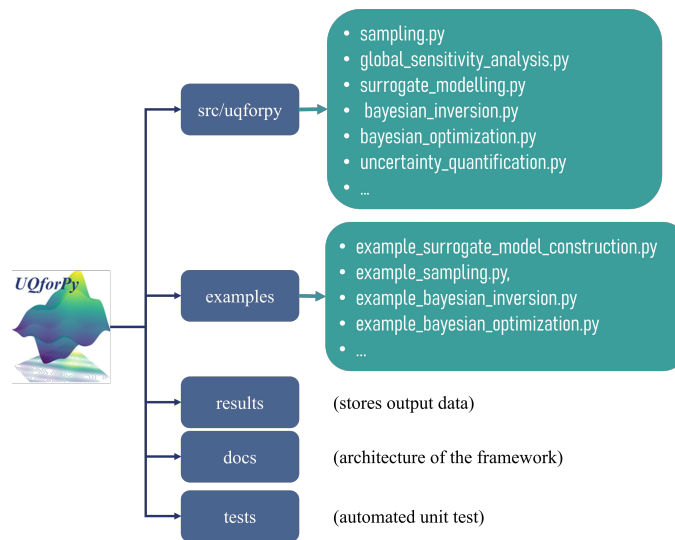


Figure C.1: Architecture of the UQforPy repository. The framework is organized into modular components supporting sampling, global sensitivity analysis, surrogate modeling, Bayesian optimization, Bayesian inversion, and uncertainty quantification.

Within `src/uqforpy`, the computational workflow is organized into two complementary

layers: (i) high-level analysis scripts, which provide user-facing interfaces for the execution of numerical workflows, and (ii) low-level algorithmic modules, which implement the underlying numerical strategies and reusable computational components.

C.1 High-level analysis scripts

The files `sampling.py`, `global_sensitivity_analysis.py`, `surrogate_modelling.py`, `bayesian_optimization.py`, `bayesian_inversion.py`, and `uncertainty_quantification.py` define the main computational workflows of the framework. These scripts orchestrate the interaction between surrogate models, sampling procedures, and inference algorithms, and are typically invoked either within user-defined applications or through the example scripts provided in the `examples` directory.

In particular:

- **`sampling.py`** implements the procedures used to generate and refine evaluation points during surrogate model construction and validation. The sampling strategies rely on the rules defined in the `sampling_rules` module.
- **`global_sensitivity_analysis.py`** performs global sensitivity analyses aimed at quantifying the influence of input parameters on the model response. The script combines sampling procedures, model evaluations, and sensitivity indicators to provide interpretable measures of parameter importance.
- **`surrogate_modelling.py`** implements the construction of surrogate models from available training data, leveraging the approximation techniques defined in the `surrogate_models` module.
- **`bayesian_optimization.py`** implements adaptive surrogate model construction strategies based on Bayesian optimization. The procedure iteratively enriches the training dataset through acquisition-driven sampling and updates the Gaussian Process surrogate model accordingly.
- **`bayesian_inversion.py`** performs parameter inference by solving inverse problems using the surrogate model. The implementation includes routines for Maximum A Posteriori (MAP) estimation, evaluation of Least Squares (LS) and Negative Least Squares (NLS) functionals, and posterior distribution analysis.
- **`uncertainty_quantification.py`** implements forward uncertainty propagation and statistical post-processing, enabling the evaluation of moments, variances, and probabilistic model responses.

C.2 Algorithmic modules

The numerical strategies underlying the analysis workflows are implemented in dedicated modules organized as subdirectories within `src/uqforpy`. These modules provide reusable algorithmic components that are invoked by the high-level scripts.

- **sampling_rules** implements sampling strategies for the generation of training and validation datasets, including uniform random sampling, Latin Hypercube Sampling (LHS), and Leja sequences. The modular design allows the integration of additional techniques such as quasi-Monte Carlo sampling, sparse-grid methods, and adaptive experimental design strategies.
- **sensitivity_indicators** provides the metrics used in global sensitivity analysis, including variance-based indicators and parameter influence measures.
- **surrogate_models** contains the implementation of surrogate modeling techniques. In the present work, Gaussian Process regression with Matérn covariance kernels is adopted as the primary surrogate model. The framework, however, supports the integration of alternative approximations, including polynomial-based methods, radial basis functions, spline interpolants, polynomial chaos expansions, and neural-network-based surrogate models.
- **bayesian_optimization_techniques** implements acquisition strategies and dataset update rules used in Bayesian optimization. These include classical acquisition functions such as Upper Confidence Bound (UCB), Expected Improvement (EI), and Probability of Improvement (PI), together with routines for adaptive sampling.
- **bayesian_inversion_techniques** provides the algorithms used in the Bayesian inversion stage, including MAP estimation procedures, evaluation of LS and NLS functionals, and local posterior approximations.
- **uncertainty_quantification_tools** implements routines for forward uncertainty propagation and statistical analysis, including moment estimation, variance decomposition, and probabilistic evaluation of model responses.

C.3 Benchmark functions and numerical utilities

The module `benchmarks` contains a collection of analytical functions commonly used for validation in optimization and inverse problem studies. The implemented benchmark suite includes the Mixed Gaussian–Periodic, Lévy, Griewank, Forrester, and Rosenbrock functions, which are employed in the numerical experiments discussed in Section 3.1.4. Additional utilities are provided in the `utils` module, including routines for numerical operations, data handling, and visualization.

C.4 Repository organization and reproducibility

The repository includes auxiliary directories supporting experimentation and reproducibility:

- **examples** provides demonstration scripts illustrating typical workflows, including surrogate modeling, Bayesian optimization, Bayesian inversion, uncertainty quantification, and sensitivity analysis.
- **tests** contains automated unit tests used to verify the correctness and robustness of the implemented algorithms.
- **docs** provides documentation describing the architecture and usage of the framework.
- **results** stores the output generated during numerical experiments, including surrogate predictions, adaptive sampling traces, inversion results, and sensitivity analysis outputs.

The repository further includes configuration files such as `requirements.txt` and `pyproject.toml`, which specify the software dependencies required to reproduce the computational experiments. The modular architecture of UQforPy facilitates integration with external high-fidelity solvers and supports future extensions of surrogate-assisted Bayesian workflows to higher-dimensional inverse problems and application-specific engineering models.

Appendix D

Adaptive Surrogate Modeling Results for One-Dimensional Benchmarks

The present appendix complements the results discussed in Section 3.1.4.1.1 by providing additional validation studies of the surrogate modeling strategy adopted in this work.

The analysis focuses on the Lévy, Griewank, and Forrester benchmark functions (Figures 3.2b–d), with the objective of systematically assessing the predictive performance of different surrogate models under uncertainty-aware adaptive sampling strategies. The study is designed to justify the selection of Gaussian Process (GP) surrogate models within the proposed framework by comparing their performance against alternative deterministic approximation techniques.

Figure D.1 reports the predictive responses obtained using different surrogate modeling approaches. The results confirm and extend the observations discussed for the Mixed Gaussian–Periodic function in Section 3.1.4.1.1.

For the Lévy function (Figure D.1a), characterized by a highly non-convex landscape with multiple local minima, GP surrogate models with both Matérn and RBF covariance kernels accurately reproduce the complex structure of the high-fidelity model. The Matérn-based surrogate model exhibits increased robustness in resolving localized irregularities and sharp transitions. In contrast, deterministic approximation methods — including Lagrangian Polynomial (LP), Legendre Expansion (LE), and Cubic Spline (CS) models — either lack sufficient flexibility to capture fine-scale variations or exhibit reduced stability in regions of strong nonlinearity.

For the Griewank function (Figure D.1b), which combines a smooth global trend with high-frequency oscillatory components, GP surrogate models again provide an accurate representation of the underlying response. Polynomial-based approximations suffer from spurious oscillations near the domain boundaries, commonly associated with Runge’s phenomenon [177], while spectral expansions fail to resolve the high-frequency content of the function.

For the Forrester function (Figure D.1c), characterized by a smooth global trend with localized nonlinear features, the GP surrogate model with Matérn kernel achieves the highest predictive accuracy. Deterministic models tend to over-smooth the response or fail to capture localized variations, leading to systematic approximation errors.

The qualitative observations are further supported by the quantitative error analysis reported in Figure D.2, where the mean squared error (MSE) is evaluated on independent validation datasets. GP surrogate models consistently achieve lower prediction errors compared to the considered deterministic alternatives, with the Matérn kernel providing the most accurate and




Figure D.1: Predictive responses of different surrogate models for the benchmark functions in Figure 3.2b–d. The solid black line represents the high-fidelity model; fuchsia, purple, and orange dotted lines correspond to LP, LE, and CS models; cyan and blue lines represent GP surrogate models with Matérn and RBF kernels. Red diamonds denote training sample locations. (a) Lévy. (b) Griewank. (c) Forrester.




Figure D.2: Mean squared error (MSE) comparison of surrogate models for the benchmark functions in Figure 3.2b–d.

stable approximation across all benchmark functions.

Overall, the results provide clear evidence that GP surrogate models, particularly when combined with Matérn covariance kernels, offer a robust and flexible approximation framework capable of capturing complex nonlinear behaviors under limited data availability. These properties motivate their adoption as the default surrogate modeling strategy within the uncertainty-aware framework proposed in this work.

Appendix E

Bayesian Inversion Results for 1D Benchmarks

The present appendix complements the results discussed in Section 3.1.4.1.1 by providing additional validation studies of the surrogate modeling strategy adopted in this work. Specifically, the parameter inference results obtained via Bayesian Inversion (BI) are reported for the Lévy and Griewank benchmark functions (Figure 3.2b–c), selected to examine the framework’s capability in the presence of localized nonlinearities and strongly multimodal landscapes.

For each benchmark, the MAP estimate, the Least Squares (LS) functional, and the corresponding Negative Least Squares (NLS) profile are analysed to characterize the structure of the inverse problem and the associated epistemic uncertainty.

Figure E.1 reports the results for the Lévy benchmark. Figure E.1(a) shows the inferred MAP estimates obtained from multiple independent optimization runs, each initialized from a different guesses point in the parameter space. All runs are based on the GP surrogate model constructed via BO (Figure 3.5b), and consistently converge to the same region.

The LS profile (Figure E.1(b)) exhibits a sharp and well-separated minimum, indicating strong local identifiability. The corresponding NLS profile (Figure E.1(c)) is sharply unimodal and concentrated around the MAP estimate $\mu = 0.76$, with standard deviation $\sigma \approx 0.02$. This concentrated posterior distribution reflects the presence of a well-posed inverse problem under the surrogate model.

Minor secondary peaks are visible in the NLS profile, but their amplitudes are negligible relative to the main mode. These features are attributed to residual numerical artifacts or minor secondary minima and do not affect the stability of the MAP solution.

The inferred posterior distribution allows for a significant reduction in the admissible parameter space, effectively isolating a narrow high-probability region. This confirms the robustness of the proposed approach in resolving inverse problems with localized nonlinear responses.

Figure E.2 shows the results for the Griewank benchmark. Figure E.2(a) presents the MAP estimates obtained from optimization runs initialized at different guesses points, using the GP surrogate constructed via BO (Figure 3.5c).

The LS profile (Figure E.2(b)) reveals a highly multimodal and periodic landscape, with multiple minima of comparable depth. As a result, the NLS profile (Figure E.2(c)) shows multiple peaks of similar amplitude, corresponding to equally plausible parameter configurations.

This inverse problem is intrinsically non-identifiable in a probabilistic sense. Despite the surrogate model’s accuracy in reconstructing the admissible solution set, the posterior distribution




Figure E.1: Parameter inference via BI results for the Lévy benchmark. (a) MAP estimates from multiple optimization runs. (b) LS functional profile. (c) NLS functional profile representing the posterior distribution approximation with high-probability region marked by the grey dotted horizontal line representing the 0.95-threshold and highlighted by the coral-coloured portion delimited by the dotted vertical lines.




Figure E.2: Parameter inference via BI results for the Griewank benchmark. (a) MAP estimates from multiple initializations. (b) LS functional profile. (c) NLS functional profile revealing non-identifiability due to periodic structure.

fails to isolate a dominant mode due to the periodic nature of the objective function.

The analyses of the Lévy and Griewank benchmarks highlight the flexibility and diagnostic power of the proposed framework. In the case of the Lévy function, the posterior distribution represented by the NLS profile is sharply concentrated and uniquely peaked, confirming the well-posedness of the inverse problem. In contrast, the Griewank case demonstrates that the framework can effectively reveal inherent non-identifiability in periodic inverse settings.

In both cases, the proposed Bayesian framework yields an interpretable characterization of the admissible parameter space, supporting robust inference and informed decision-making even in the presence of structural ambiguity.

Appendix F

Supplementary Surrogate-Based Posterior Analysis

The present appendix provides a supplementary posterior analysis based on surrogate-assisted Markov Chain Monte Carlo (MCMC) sampling [105, 158, 159]. The analysis is intended as a complementary validation tool rather than as a methodological component of the proposed framework, and aims to illustrate regimes in which sampling-based inference becomes necessary to accurately characterize the posterior distribution.

The study is conducted using the Gaussian Process (GP) surrogate model constructed in the surrogate modeling phase, and focuses on the one-dimensional Mixed Gaussian–Periodic benchmark function (Figure 3.2a), selected as a representative case of a strongly non-convex inverse problem with expected multimodal posterior structure.

The objective is to obtain an independent representation of the posterior distribution in a setting where local Gaussian approximations around a single maximizer are not theoretically justified. In such regimes, the posterior distribution may exhibit multiple disconnected regions of high probability, and sampling-based exploration provides a direct means of resolving its global structure.

The marginal posterior densities estimated from independent MCMC chains are reported in Figure F.1. Each chain is initialized from a different starting point in the parameter space in order to assess convergence behavior and sensitivity to initialization.

The results consistently reveal a multimodal posterior structure, characterized by a dominant region of high probability together with additional secondary modes. The dominant mode is robustly identified across most chains, while the exploration of secondary modes exhibits variability depending on the initialization and stochastic trajectory of each chain. This behavior reflects the intrinsic non-convexity of the inverse problem and the presence of multiple admissible parameter configurations.

A global visualization of the posterior structure is provided in Figure F.2, where the marginal posterior densities obtained from all chains are superimposed. The overlay highlights the consistency of the dominant mode and the variability associated with secondary modes, providing a compact representation of posterior complexity.

The observed multimodality confirms that the assumptions underlying local Laplace approximations around a single MAP estimate are not satisfied for the Mixed Gaussian–Periodic benchmark. In particular, the posterior distribution cannot be adequately represented by a single Gaussian approximation centered at a dominant mode.

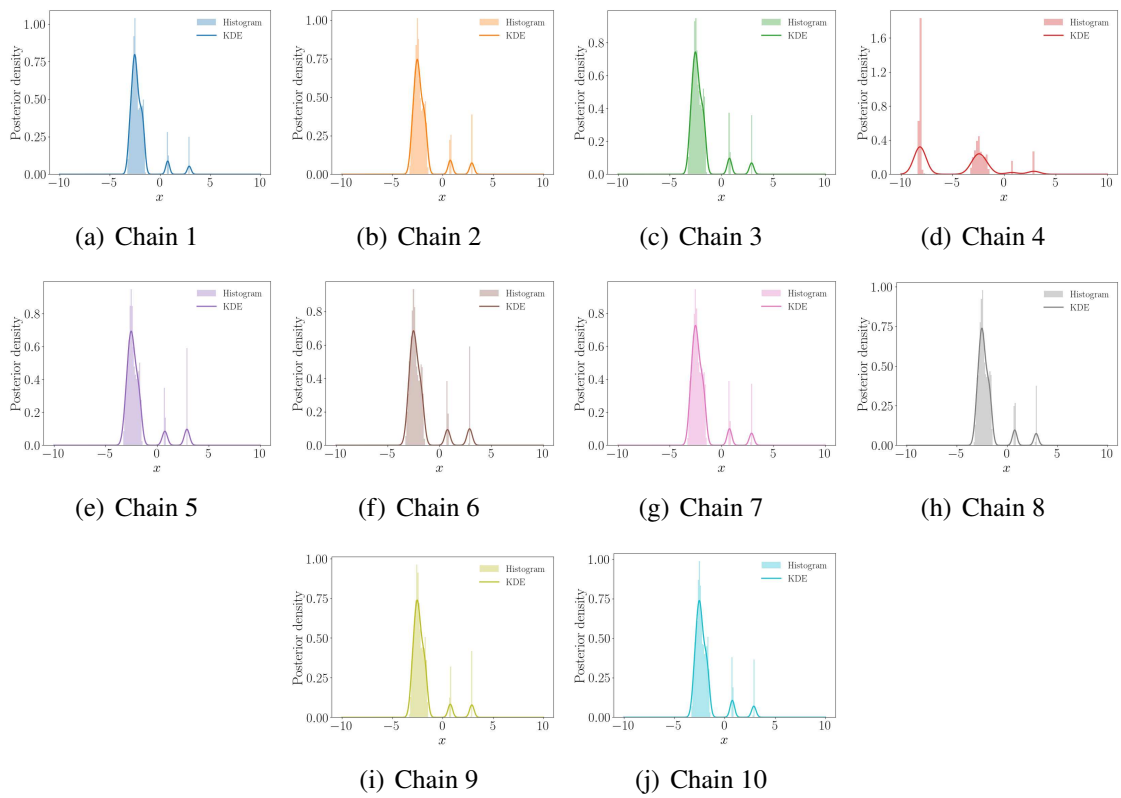


Figure F.1: Marginal posterior densities obtained from surrogate-assisted MCMC sampling for the Mixed Gaussian–Periodic benchmark. For each chain, the empirical histogram of the samples is shown together with a kernel density estimate.




Figure F.2: Superimposed marginal posterior density estimates obtained from all surrogate-assisted MCMC chains. The dominant mode is consistently identified, while secondary modes are sampled with varying frequency.

For this reason, the present work does not enforce a local Gaussian approximation in such regimes. Instead, the negative least-squares (NLS) functional is employed as a computationally inexpensive diagnostic tool to assess the structure of the posterior distribution and to identify regions of high posterior probability. The NLS profile provides a reliable indication of the location and relative prominence of posterior modes, while avoiding the need for expensive sampling-based exploration.

For completeness, a reference posterior density is also constructed through a dense grid-based evaluation of the surrogate-induced likelihood over the admissible parameter domain. The resulting posterior is shown in Figure F.3, obtained by normalizing the likelihood under a uniform prior.

The grid-based posterior clearly exhibits multiple well-separated modes and is fully consistent with the LS and NLS profiles discussed in Section 3.1.3.3. This deterministic representation provides a global reference against which the MCMC results can be qualitatively assessed.

From a methodological perspective, the results highlight the complementary roles of different inference strategies. Local Laplace approximations are appropriate when the posterior distribution is effectively unimodal and the inverse problem is well posed. In contrast, sampling-based approaches such as MCMC are required to accurately resolve multimodal posterior structures and to capture the full extent of epistemic uncertainty.

The NLS-based analysis adopted in the present work provides an intermediate diagnostic tool that enables the identification of high-probability regions at negligible computational cost, while preserving consistency with the underlying probabilistic formulation. When multimodality is detected, as in the present benchmark, MCMC-based exploration can be employed as a complementary validation strategy.




Figure F.3: Grid-based posterior density for the Mixed Gaussian–Periodic benchmark, obtained by evaluating the surrogate-induced likelihood over a dense discretization of the parameter space. The posterior exhibits multiple well-separated modes.

Overall, the results support a principled and adaptive use of inference techniques, in which Laplace approximations, NLS-based diagnostics, and sampling-based methods are employed according to the structural properties of the posterior distribution, rather than as interchangeable tools.