SISSA
Scuola Internazionale Superiore di Studi Avanzati

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

Finanziato dall'Unione Europea – *Next Generation EU*" - Missione 4, Componente 2 del Piano Nazionale di Ripresa e Resilienza (PNRR), Investimento 3.3 "Dottorati innovativi che rispondono ai fabbisogni di innovazione delle imprese" - CUP G41J22000270008

Mathematics Area – PhD course in
Mathematical Analysis, Modelling, and Applications

# Parameterized Orthotropic Plate Systems: Numerical Methods and Model Reduction

*From High-Fidelity $C^0$ Interior-Penalty Discretizations to Reduced-Order Models for Plates on Elastic Foundations*

Candidate:
Anantha Krishnan Orunnukaran Mani

Advisors:
Prof. Giovanni Noselli
Dr. Federico Pichi
Prof. Annalisa Quaini
Prof. Gianluigi Rozza

Academic Year 2025–26



**SISSA – International School for Advanced Studies**

Mathematical Analysis, Modelling, and Applications

XXXVIII cycle

The present dissertation, submitted by Anantha Krishnan Orunnukaran Mani under the supervision of Prof. Giovanni Noselli, Dr. Federico Pichi, Prof. Gianluigi Rozza (SISSA, Italy), and of Prof. Annalisa Quaini (University of Houston, USA), is submitted in fulfillment of the requirements for the degree of *Philosophiæ Doctor* in Mathematical Analysis, Modelling, and Applications at SISSA – International School for Advanced Studies. According to Article 1, paragraph 4, of the SISSA Statute, published in G.U. No. 36 on 13 February 2012, the aforementioned degree is equivalent to the title of Research Doctorate in Mathematics.

Final Examination: September 23, 2026

Reviewers:

Prof. Davide Bigoni, University of Trento, Italy
Prof. Andrea Mola, IMT School for Advanced Studies Lucca, Italy

Board of Examiners:

Prof. Francesco Dal Corso, University of Trento, Italy
Prof. Andrea Mola, IMT School for Advanced Studies Lucca, Italy
Prof. Michele Girfoglio, Università degli Studi di Palermo, Italy
Dr. Marco Tezzele, Emory University, USA
Prof. Giovanni Noselli, SISSA – International School for Advanced Studies, Italy
Prof. Gianluigi Rozza, SISSA – International School for Advanced Studies, Italy

Thanks be to science!





Trieste, Italia

# Acknowledgments

My advisor, **Prof. Gianluigi Rozza**, has shaped both this dissertation and the way I approach research and scientific questions. I am deeply grateful for his guidance throughout this work, for the trust and encouragement he gave me, and for the scientific perspective he brought to it.

Discussions with **Prof. Giovanni Noselli** repeatedly returned to the mechanics behind the equations: what a model represents, what an assumption costs, and what a numerical result allows us to conclude. I thank him for keeping these questions in view. With **Dr. Federico Pichi**, our discussions of the numerical and reduced-order work made many parts of this thesis clearer and stronger. I am particularly grateful for his patience, his close reading, and the care he took in guiding me through that work. In my conversations with **Prof. Annalisa Quaini** of the University of Houston, USA, I valued her thoughtful questions and generous advice. I thank her warmly for these exchanges and for the clarity she consistently brought to them.

I also thank **Dr. Luca Tosolini**, **Dr. Luca Giubilato**, **Dr. Alessandro Corso**, and **Andrea Martini** for the industrial support and perspective they brought to this doctoral work through the PNRR-supported SISSA–EPS collaboration. Andrea has been involved from the beginning, and I am especially grateful for his help in connecting the academic and industrial sides of the project throughout.

Some thanks are harder to express. I am deeply grateful to my father, **Mani O. K**, for everything he gave me long before I began this PhD. I write his name here knowing that these few lines can only express part of that gratitude. To my mother, **Suseela**, thank you for your strength, your love, and your faith in me through all these years and across the distance.

To **Gopika**, my **Poppy**, thank you for sharing these years with me, for your patience, and for all the laughter and love along the way. On the longest days, you reminded me how much there was to life beyond the work on my desk.

Friendships formed in very different places have sustained me before and during this PhD. To my friends from **SVNCS**, **Carmel**, **SRM**, **Sapienza**, and **SISSA**, thank you for staying close through the years, across cities and countries, and as our lives changed. Our conversations, jokes, meals, arguments, reunions, and ordinary moments together have meant more to me than I can fully express here. And finally, **Trieste**. Most of this thesis was written here, and the city became part of the memory of the work itself: the bora, the sea, the long walks, and the views I learned to recognize without thinking. I will carry those details with me when I think back on these years.

## Funding and Institutional Support

This doctoral research was supported by a doctoral fellowship funded through the Italian National Recovery and Resilience Plan (PNRR), financed by the European Union through *Next Generation EU*, Mission 4, Component 2, Investment 3.3, "Dottorati innovativi che rispondono ai fabbisogni di innovazione delle imprese," and administered by the Italian Ministry of University and Research (MUR), CUP G41J22000270008.

I gratefully acknowledge EPS Italia Srl Società Benefit for hosting the six-month industrial research period carried out within the SISSA–EPS Italia PhD project and for the technical collaboration that accompanied this work. I also acknowledge the SISSA mathLab group for supporting my participation in scientific conferences and advanced training schools.

Finanziato dall'Unione europea NextGenerationEU | Ministero dell'Università e della Ricerca | Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

Finanziato dall'Unione Europea – *Next Generation EU*" - Missione 4, Componente 2 del Piano Nazionale di Ripresa e Resilienza (PNRR), Investimento 3.3 "Dottorati innovativi che rispondono ai fabbisogni di innovazione delle imprese" - CUP G41J22000270008

**Doctoral thesis**

August 2026

# Abstract

This dissertation develops and evaluates a unified high-fidelity and reduced-order framework for parameterized thin orthotropic Kirchhoff–Love plates. Motivated primarily by monolithic and panelized flooring systems on compliant foundations, the formulations also apply to plates with parameter-dependent material, support, loading, thermal, and dynamic properties. Three connected model classes are considered: static mechanics with unilateral plate–foundation interaction, steady thermomechanical bending with displacement-dependent thermal exchange, and transient dynamics. The midsurface may be partitioned into non-overlapping structural panels, with transverse displacement and normal rotation transmitted weakly across physical seams.

The static mechanical model combines aligned orthotropic bending, parameterized loads and rigidities, panel transmission, and a regularized unilateral Winkler law distinguishing compression from uplift. The fourth-order problem is approximated by a symmetric $C^0$ interior-penalty Galerkin method on $H^1$-conforming Lagrange spaces. Numerical stabilization on element facets is kept distinct from the displacement and normal-rotation penalties across physical seams. An active-set Newton method resolves the unilateral contribution using an algebraic residual stopping criterion.

The thermomechanical extension couples the plate to steady three-dimensional heat conduction with exposed-face convection, long-wave radiation, absorbed solar flux, and substrate exchange. A centered first moment of the through-thickness temperature increment defines the plate thermal driver $\theta$, measured in $\mathrm{K\,m^{-1}}$, and the associated orthotropic thermal bending moments. Displacement modifies lower-face heat transfer through a contact–gap exchange law, while the resulting temperature and thermal driver modify mechanical equilibrium. A relaxed partitioned fixed-point iteration alternates the thermal solve, thermal-to-plate reduction, nonlinear plate solve, and displacement relaxation.

The transient extension adds transverse translational inertia, initial displacement and velocity, time-dependent loading, optional algebraic Rayleigh damping, and a consistent mass matrix. A common active-degree restriction removes zero-mass entries in panelized systems, and the average-acceleration Newmark method advances the resulting second-order system. Each trajectory is computed with one prescribed effective Winkler coefficient, so the transient study represents a fixed foundation regime with transverse translational inertia only.

The full-order models (FOMs) are assessed through complementary numerical verification. The mechanical approximation is compared with a truncated Navier-series reference and an independent three-dimensional solid model, examined through dual-reference spatial refinement, and evaluated against monolithic references for two-, three-, four-, and six-panel layouts. The thermomechanical solver is assessed through an auxiliary analytical comparison, plate-only and coupled refinement, partitioned-iteration histories, and a cost decomposition showing that the three-dimensional thermal stage and thermal-to-plate reduction dominate recurrent coupled-solve time. The transient implementation is checked through analytical and finite element modal comparisons, time-domain references, direct-to-modal agreement, undamped energy conservation, harmonic and localized-load responses, and panel-to-monolithic consistency. Together, these studies provide numerical verification for the tested configurations, discretizations, and parameter values.

Because repeated high-fidelity evaluation is expensive, metric-aware Proper Orthogonal Decomposition (POD) provides a common compression framework for reduced-order models (ROMs). For static mechanics, a non-hyper-reduced intrusive POD–Galerkin model is compared with interpolation and regression of POD coordinates. The thermomechanical FOM remains physically coupled, while its ROM assessment uses two independent one-field workflows generated from the same accepted coupled states. The displacement and thermal-driver fields retain separate output spaces, metrics, training-only bases, predictors, units, and error measures while using matched sample assignments within each parameterization. The non-intrusive methods comprise radial basis function and linear POD interpolation (PODI–RBF and PODI–Linear), Gaussian-process and neural-network regression of POD coordinates (POD–GPR and POD–NN), a POD-enhanced deep-learning ROM (POD–DL-ROM), and direct deep-learning ROMs (DL-ROMs).

For transient dynamics, displacement is the only reduced output. Complete trajectories are assigned to disjoint training, validation, and test subsets before flattening, and only training trajectories enter POD construction, preprocessing, and model fitting. PODI–RBF, PODI–Linear, POD–GPR, POD–NN, and POD–DL-ROM map physical parameters and query time to displacement coordinates. Time is an independent query coordinate, yielding a time-conditioned, non-causal displacement-field predictor.

Across the sampled parameter domains, the displacement and thermal-driver fields exhibit strong low-dimensional structure, although predictive performance depends on the complexity and coverage of the learned parameter-to-output map. In the mechanical study, POD–GPR achieves accuracy comparable to intrusive POD–Galerkin at substantially lower online cost, while the acceleration of the intrusive model remains limited by full-order residual and Jacobian assembly. In the thermomechanical study, POD–GPR gives the lowest mean thermal-driver error across all parameterizations and the lowest displacement error when the response is parameterized by mechanical loading and absorbed solar flux. Direct DL-ROM gives the best displacement predictions when the response is governed by foundation stiffness and boundary temperatures and in the broader multiparameter setting. In transient dynamics, the load-amplitude maps are predicted close to the projection benchmark, whereas stiffness-dependent response is more demanding because changes in modal frequencies and accumulated phase increase the error of the parameter–time coordinate map. Direct POD projection (POD–Proj) is used as a non-predictive compression diagnostic, and transient online timing is reported as a complete-trajectory-to-single-state cost ratio.

Overall, the dissertation establishes a coherent hierarchy from parameterized plate modeling and verified $C^0$ interior-penalty approximation to intrusive, POD-based non-intrusive, and deep-learning ROMs for many-query prediction.

**Keywords:** *Orthotropic Kirchhoff–Love plates; unilateral Winkler foundation; panel interfaces; $C^0$ interior-penalty Galerkin method; thermomechanical coupling; transient dynamics; Proper Orthogonal Decomposition; reduced-order modeling; many-query prediction.*

# Contents

**Abstract** vi

**Contents** viii

**Foundations: Motivation and Theoretical Framework** 1

**1 Introduction** 2
1.1 Motivation and thesis perspective . . . 2
1.2 Common modeling framework and three physical studies . . . 3
1.3 High-fidelity discretization and verification . . . 4
1.4 Many-query reduction and reduced-order prediction . . . 5
1.5 Dissertation scope and contributions . . . 7
1.6 Organization of the thesis . . . 8

**2 Theoretical Background for Parameterized Plate Models** 10
2.1 Role and scope of the background chapter . . . 10
2.2 Thin-plate idealization and midsurface variables . . . 10
2.3 Kirchhoff–Love kinematics and bending measures . . . 11
2.4 Aligned orthotropic bending and energetic structure . . . 12
2.5 Compliant supports and panel interfaces . . . 12
2.6 Variational structure and boundary conditions . . . 13
2.7 Discretization of the fourth-order plate operator . . . 14
2.8 Thermomechanical coupling at plate level . . . 15
2.9 Transient dynamics and modal structure . . . 16
2.10 Parameterized states, outputs, and model reduction . . . 17
2.11 Summary and handoff to the study chapters . . . 18

**I Static Mechanics of Orthotropic Plates** 20

**3 Parameterized Mechanical Kirchhoff–Love Plate Model** 21
3.1 Role of the mechanical model in the thesis . . . 21
3.2 Geometric setting and panel partition . . . 21
3.3 Kirchhoff–Love kinematics and sign conventions . . . 22
3.4 Orthotropic bending resultants . . . 23
3.5 Unilateral Winkler foundation and uplift regularization . . . 23
3.6 Interface transmission across panel partitions . . . 24
3.7 Energetic formulation . . . 24
3.8 Weak formulation . . . 25
3.9 Parameter dependence and nonlinear structure . . . 26
3.10 Summary of the mechanical model . . . 27

**4 Full-Order Approximation of the Mechanical Plate Model** 28
4.1 Role of the high-fidelity approximation . . . 28
4.2 Discrete spaces and panel-wise finite element meshes . . . 28
4.3 $C^0$ interior-penalty discretization inside panels . . . 29
4.4 Weak enforcement of panel-interface continuity . . . 30
4.5 Discrete unilateral foundation term . . . 31
4.6 Mechanical full-order residual problem . . . 31

4.7 Nonlinear solution strategy . . . 32
4.8 Verification against analytical and three-dimensional references . . . 33
4.9 Mesh-refinement convergence study . . . 37
4.10 Penalty-coupled panel-partition assessment . . . 39
4.11 Summary of the full-order model . . . 41

**5 Reduced-Order Modeling of the Mechanical Plate Problem 42**
5.1 Many-query setting and reduced-order objective . . . 42
5.2 POD basis construction and intrusive POD–Galerkin reduction . . . 42
5.3 Non-intrusive POD-coordinate surrogates . . . 44
5.4 Assessment protocol and offline–online accounting . . . 48
5.5 ROM assessment for high-dimensional mechanical parametrization . . . 49
5.6 ROM assessment for penalty-coupled panel systems . . . 52
5.7 ROM assessment for multi-panel configurations . . . 56
5.8 Summary and limitations of the mechanical ROMs . . . 59

**II Thermomechanics and Transient Dynamics of Orthotropic Plates 61**

**6 Thermomechanical Kirchhoff–Love Plate Formulation 62**
6.1 Thermomechanical extension of the mechanical model . . . 62
6.2 Geometry, panel partition, and three-dimensional thermal domain . . . 62
6.3 Thermal subproblem . . . 64
6.4 Weak formulation of the thermal problem . . . 65
6.5 Thermal-to-plate reduction . . . 66
6.6 Thermal bending moments . . . 66
6.7 Mechanical equilibrium with thermal loading . . . 68
6.8 Partitioned staggered thermomechanical coupling . . . 70
6.9 Parameter dependence and feedback structure . . . 72
6.10 Summary of the thermomechanical formulation . . . 72

**7 Full-Order Thermomechanical Approximation and Coupled-Solver Assessment 74**
7.1 Role and scope of the thermomechanical high-fidelity model . . . 74
7.2 Finite element approximation of the thermal subproblem . . . 74
7.3 Discrete reconstruction of the plate thermal driver . . . 75
7.4 $C^0$ interior-penalty plate approximation with thermal bending . . . 76
7.5 Partitioned thermomechanical full-order solves . . . 78
7.6 Navier-series reference for the simply supported auxiliary problem . . . 79
7.7 Verification and convergence studies . . . 79
7.8 Summary of the thermomechanical full-order model . . . 86

**8 Reduced-Order Modeling of the Thermomechanical Plate Problem 87**
8.1 Thermomechanical many-query setting and reduction objective . . . 87
8.2 Reduced output maps and snapshot databases . . . 87
8.3 Metric-aware field-wise POD spaces . . . 89
8.4 Hybrid projection-based thermomechanical ROM . . . 90
8.5 Non-intrusive one-field ROMs . . . 91
8.6 Error measures and assessment protocol . . . 93
8.7 Numerical assessment . . . 94
8.8 Interpretation, limitations, and chapter conclusions . . . 100

**9 Parameterized Transient Dynamics of Orthotropic Kirchhoff–Love Plates 102**
9.1 Role and scope of the transient model . . . 102
9.2 Geometry, panel partition, and notation . . . 102
9.3 Kirchhoff–Love kinematics and orthotropic bending . . . 104

9.4 Kinetic, elastic, foundation, interface, and loading contributions . . . . . 105
9.5 Hamilton's principle and weak evolution problem . . . . . 107
9.6 Strong form, boundary conditions, and initial data . . . . . 108
9.7 Parameterized transient solution map . . . . . 109
9.8 Summary and handoff to the full-order approximation . . . . . 110

**10 Full-Order Approximation and Verification of the Transient Plate Model 111**
10.1 Role and scope of the transient full-order model . . . . . 111
10.2 Panel-aligned spatial discretization . . . . . 111
10.3 Semidiscrete dynamics and active degrees of freedom . . . . . 114
10.4 Time integration and snapshot storage . . . . . 115
10.5 Analytical and modal references . . . . . 116
10.6 Verification framework and results . . . . . 120
10.7 Summary and handoff to the reduced-order model . . . . . 126

**11 Non-Intrusive Reduced-Order Models for Parameterized Transient Dynamics 127**
11.1 Many-query objective and assessment structure . . . . . 127
11.2 Time-dependent displacement output . . . . . 127
11.3 Snapshot campaigns and complete-trajectory data split . . . . . 128
11.4 Training-only POD construction and projection diagnostic . . . . . 130
11.5 Active non-intrusive coordinate models . . . . . 131
11.6 Data partition, error measures, and timing protocol . . . . . 133
11.7 Reduced-order results . . . . . 134
11.8 Cross-campaign interpretation and conclusions . . . . . 140

**Conclusions and Future Outlook 141**

**A Computational Implementation and Reproducibility 146**
A.1 Purpose and scope . . . . . 146
A.2 Software, artifacts, and shared finite element architecture . . . . . 146
A.3 Study-specific full-order workflows . . . . . 148
A.4 Snapshot, output-transfer, and reduced-order pipelines . . . . . 150
A.5 Reduced-order model architectures and hyperparameters . . . . . 152
A.6 Code and computational reproducibility . . . . . 156

**Bibliography 157**

# List of Figures


1.1 Representative examples of panelized flooring systems in large-event settings . . . . . 3
1.2 Unified model–FOM–ROM architecture of the dissertation . . . . . 7

2.1 Thin-plate idealization and Kirchhoff–Love kinematics . . . . . 11

3.1 Geometric setting for the mechanical Kirchhoff–Love plate model . . . . . 22

4.1 Verification summary for the monolithic simply supported plate . . . . . 36
4.2 Moment fields for the monolithic full-order solution . . . . . 37
4.3 Relative displacement-error convergence for the mechanical FOM . . . . . 39
4.4 Panel-partition layouts for the orthotropic benchmark . . . . . 40
4.5 Absolute error fields for penalty-coupled panel partitions . . . . . 40
4.6 Relative error for panel-partition configurations . . . . . 41

5.1 Offline sampling and POD spectrum for the high-dimensional mechanical ROM test . 50
5.2 Validation-error sensitivity to the retained POD dimension . . . . . 50
5.3 Test-space accuracy for the high-dimensional mechanical ROM test . . . . . 51
5.4 Representative field comparison for the high-dimensional mechanical ROM test . . . . 52
5.5 Offline sampling and POD spectrum for the two-panel ROM test . . . . . 53
5.6 Deflection profiles for the two-panel ROM test . . . . . 54
5.7 Representative field comparison for the two-panel ROM test . . . . . 55
5.8 Test-space accuracy for the two-panel ROM test . . . . . 55
5.9 Deflection profiles for the multi-panel ROM tests . . . . . 57
5.10 Absolute error fields for the multi-panel ROM tests . . . . . 58
5.11 Test-space accuracy for the multi-panel ROM tests . . . . . 59

6.1 Thermomechanical plate geometry and panel partition . . . . . 64

7.1 Auxiliary Navier-series comparison for the thermomechanical FOM . . . . . 81
7.2 Plate-only refinement for the thermomechanical FOM . . . . . 82
7.3 Coupled thermomechanical refinement and iteration convergence . . . . . 84

8.1 Displacement and thermal-driver POD spectra for the thermomechanical ROMs . . . . 94
8.2 Reduced-dimension sensitivity for the Case 2 thermomechanical ROMs . . . . . 96
8.3 Held-out accuracy of the thermomechanical one-field ROMs . . . . . 97
8.4 Online timing ratios for the thermomechanical one-field ROMs . . . . . 99
8.5 POD projection displacement and error fields for thermomechanical ROMs . . . . . 99
8.6 Predictive error fields for thermomechanical one-field ROMs . . . . . 100

9.1 Transient plate geometry and panel partition . . . . . 103

10.1 Natural-frequency verification for the transient FOM . . . . . 121
10.2 Time-domain verification for the transient FOM . . . . . 123
10.3 Energy and harmonic-response verification for the transient FOM . . . . . 124
10.4 Localized-load and panel-consistency verification for the transient FOM . . . . . 125

11.1 POD spectra for the transient ROM campaigns . . . . . 131
11.2 Reduced-dimension sensitivity for the Case 2 transient ROM . . . . . 135
11.3 Held-out accuracy of the transient ROM campaigns . . . . . 135
11.4 POD projection fields for the transient ROM campaigns . . . . . 137
11.5 Predictive error fields for the transient ROM campaigns . . . . . 137
11.6 Held-out transient probe histories and absolute errors . . . . . 139

A.1 Common POD projection, interpolation, and regression architecture . . . . . 153
A.2 Implemented POD–NN architecture . . . . . 154
A.3 Implemented POD–DL-ROM and direct DL-ROM architectures . . . . . 155

# List of Tables

4.1 Monolithic verification setup for the mechanical FOM . . . . . 34
4.2 Displacement comparison for the monolithic verification case . . . . . 35

5.1 Average ROM performance for the high-dimensional mechanical test . . . . . 51
5.2 Average ROM performance for the two-panel mechanical test . . . . . 55
5.3 Average ROM performance for the multi-panel mechanical tests . . . . . 59

7.1 Plate-only convergence rates for the thermomechanical FOM . . . . . 82
7.2 Coupled convergence rates for the thermomechanical FOM . . . . . 84

8.1 Parameter domains for the thermomechanical ROM cases . . . . . 88
8.2 Data partitions and output-space dimensions for the thermomechanical ROMs . . . . 89
8.3 Non-intrusive thermomechanical one-field ROM configurations . . . . . 93
8.4 Held-out accuracy of the thermomechanical one-field ROMs . . . . . 98

10.1 Material parameters for transient FOM verification . . . . . 120
10.2 Natural-frequency verification summary for the transient FOM . . . . . 121
10.3 First eight benchmark frequencies for the transient FOM . . . . . 122
10.4 Verification summary for the transient FOM . . . . . 125

11.1 Full-order settings for the transient ROM campaigns . . . . . 129
11.2 Data partitions and POD dimensions for the transient ROM campaigns . . . . . 129
11.3 Predictive models for the transient ROM campaigns . . . . . 132
11.4 Held-out accuracy of the transient ROM campaigns . . . . . 136

A.1 Study-specific computational states and online operations . . . . . 147
A.2 Principal ROM architectures and implementation . . . . . 154

**Foundations**

# Motivation and Theoretical Framework

> These opening chapters establish the engineering motivation, scope, and shared mathematical language of the dissertation. They introduce monolithic and panel-partitioned orthotropic Kirchhoff–Love plates on elastic foundations and connect their variational structure to $C^0$ interior-penalty approximation, thermomechanical and transient extensions, and model reduction. Together, they provide the *conceptual and theoretical foundations* needed to read the three model–FOM–ROM hierarchies while reserving *study-specific formulations, algorithms, and evidence* for the chapters that follow.

# Introduction 1

1.1 Motivation and thesis perspective . . . . . 2
1.2 Common modeling framework and three physical studies . . . . . 3
1.3 High-fidelity discretization and verification . . . . . 4
1.4 Many-query reduction and reduced-order prediction . 5
1.5 Dissertation scope and contributions . . . . . 7
1.6 Organization of the thesis 8

## 1.1 Motivation and thesis perspective

Thin plate-like structures occur in engineering systems for which slenderness, lightweight construction, and geometric modularity are central design features. Panelized temporary flooring is a representative example. Such systems must carry spatially heterogeneous loads, distribute them over compliant and sometimes sensitive substrates, and remain practical to transport, install, remove, and reuse [1, 2]. Their assessment is therefore not exhausted by a single nominal calculation. Material properties, support conditions, seam behavior, load locations and amplitudes, and environmental data may all vary from one deployment scenario to another. Predictive simulation consequently becomes a many-query task supporting verification, sensitivity analysis, design exploration, and rapid evaluation under changing operating conditions [3, 4].

The application setting of this dissertation is also industrially motivated. The work is informed by collaboration with EPS Italia Srl on temporary pedestrian and drivable event-flooring systems. EPS product documentation describes modular flooring used in stadiums, arenas, festival grounds, temporary access routes, and large venues, where panels may support pedestrians, service vehicles, heavy equipment, stages, grandstands, or other temporary infrastructure [5–8]. These systems pose questions that are not represented adequately by an isolated, perfectly supported single plate. Their response depends on directional panel stiffness, load transfer through mechanical seams, interaction with the underlying support, and, in exposed environments, thermal actions. Reducing unnecessary conservatism is relevant not only to deflection control, but also to material use, handling effort, transport, reusability, and deployment logistics. The broader literature on portable matting and temporary-road systems similarly identifies removability and repeated use as defining features of this class of infrastructure [2].

Representative deployable flooring systems and installation settings are shown in Figure 1.1.

For sufficiently slender plates whose transverse-shear deformation is not dominant, Kirchhoff–Love theory provides a natural midsurface description of flexure [11–14]. It retains the leading bending and twisting mechanisms on the midsurface and is compact enough to support variational analysis, finite element approximation, and model reduction. The practical systems considered here nevertheless depart from the simplest classical setting in several important ways. Their effective stiffness may be orthotropic; the support may act differently in compression and uplift; the global layout may consist of multiple panels with mechanically significant seams; thermal gradients may produce bending moments; and time-dependent loading may excite vibration and transient response. These additions interact with the numerical approximation and with the definition of the field that is later reduced.

The principal challenge is therefore not merely to produce a fast predictor. It is to construct three parameterized model–FOM–ROM hierarchies in which the governing assumptions, high-fidelity discretization, numerical verification, output map, and reduced-order approximation




**Figure 1.1:** Representative examples of deployable panelized flooring systems in large-event settings. Top: modular drivable flooring transported and deployed inside an arena. Bottom left: temporary access panels supporting heavy equipment near a large scaffold structure. Bottom right: lightweight panel handling and placement in a stadium. The examples illustrate the modular structural systems that motivate the plate-based modeling framework of this dissertation. Image sources: [8–10].

remain mutually consistent. The mechanical study develops nonlinear static mechanics for orthotropic Kirchhoff–Love plates on a regularized unilateral Winkler foundation. The thermomechanical study couples a three-dimensional steady heat problem to plate bending through a centered through-thickness thermal driver. The transient-dynamics study treats prescribed-regime transverse response for monolithic and panel-partitioned plates. Each hierarchy terminates in an output-specific reduced-order assessment supported by controlled data partitions and held-out predictive evidence.

## 1.2 Common modeling framework and three physical studies

The three physical studies share a common structural idealization. The plate is represented by its midsurface, its orthotropy directions are aligned with the in-plane Cartesian axes, and the primary structural unknown is the transverse displacement. Panelized layouts are described by non-overlapping subdomains whose traces remain independent until interface terms are applied. The thesis uses an upward-positive convention for both displacement and transverse loading, so a downward load has negative sign. The exact resultants, spaces, signs, and study-specific assumptions are fixed in the corresponding model chapters; Chapter 2 establishes the common plate, interface, variational, and reduction concepts used by all three studies.

The *mechanical study* considers static bending under mechanical loading. Aligned orthotropic rigidities determine the bending and twisting response, while a local Winkler law represents the substrate [15]. To distinguish compression from uplift, the active foundation stiffness depends on the sign of the displacement, following the compression-only foundation setting [16]. A small positive uplift-side stiffness is introduced

here to regularize complete loss of contact. The problem is linear on a fixed compression–uplift partition but nonlinear overall because that partition depends on the unknown displacement. For a multi-panel assembly, transverse displacement and normal rotation are transmitted weakly across the physical seams.

The *thermomechanical study* preserves this static mechanical structure and adds a steady temperature field on the three-dimensional plate body. Conduction in the solid is combined with exposed-face convection and long-wave radiation, absorbed solar flux, and lower-face thermal exchange with the substrate [17–20]. A centered first through-thickness moment of the temperature increment defines the two-dimensional plate thermal driver $\theta$, with units $\mathrm{K\,m^{-1}}$; $\theta$ is not temperature. It generates the orthotropic thermal bending moments used by the plate problem. The primary formulation is physically coupled: the current displacement modifies the lower-face contact–gap heat exchange, while the resulting thermal field and thermal driver modify mechanical equilibrium. A relaxed partitioned fixed-point iteration alternates the three-dimensional thermal solve, thermal-to-plate reduction, nonlinear plate solve, and displacement relaxation [21]. The one-way evaluation is the single-pass specialization of this coupled algorithm.

The *transient-dynamics study* extends the common plate framework in a different direction. It adds transverse translational inertia, initial displacement and velocity, time-dependent loading, and optional algebraic Rayleigh damping after spatial discretization. One effective Winkler coefficient is selected before a trajectory is advanced and remains fixed over the time interval. For every selected support regime and physical parameter instance, the evolution problem is therefore linear. It supports free-vibration, harmonic-response, and transient analyses, but does not model time-dependent detachment, recontact, impact, or compression–uplift switching. The governing model retains transverse translational inertia only; rotary inertia, transverse-shear deformation, in-plane dynamics, geometric nonlinearity, and thermomechanical coupling remain outside its scope. The transient campaigns used for model reduction are undamped.

The three model classes reuse the same plate, support, panel, discretization, and many-query vocabulary while retaining distinct state equations, physical assumptions, and outputs. This common structure supports cumulative exposition and cross-study comparison without merging the physical states or reduced outputs.

## 1.3 High-fidelity discretization and verification

Kirchhoff–Love plate equations are fourth order in space. A conforming Galerkin approximation would therefore require globally $C^1$-continuous trial functions, which are restrictive on general two-dimensional meshes and inconvenient for panel-partitioned geometries [22–24]. The high-fidelity plate solvers in this dissertation instead use symmetric $C^0$ interior-penalty Galerkin ($C^0$-IPG) formulations on standard $H^1$-conforming Lagrange spaces [25–27].

A thesis-wide distinction is maintained between two types of internal boundary. Mesh facets lying strictly inside a physical panel are *numerical* interfaces: consistency, symmetry, and stabilization terms there recover the regularity missing from the $C^0$ approximation. A seam between

two physical panels is a *mechanical* interface: separate penalty terms transmit displacement and normal rotation between independent panel fields. Conflating these roles would obscure both the physical model and its numerical implementation. Related weak-coupling strategies for higher-order plate and shell discretizations include penalty, mortar, Nitsche-type, and super-penalty formulations [28–32]; the present work adopts a panel-aligned penalty construction tailored to the common finite element architecture.

The nonlinear, coupled, and dynamic algorithms follow the structure of their parent models. The static unilateral problem is solved by an active-set Newton method with an algebraic residual stopping test [33]. In the thermomechanical problem, Newton solves resolve the thermal and mechanical subproblems, while a relaxed partitioned iteration transmits displacement and thermal information between them and monitors finite element $L^2(\Omega)$ updates of both displacement and thermal driver. In the dynamic problem, a consistent mass matrix and the plate stiffness define a second-order system advanced by the average-acceleration Newmark method [34]. The panel-wise mixed realization requires a common restriction to the dynamically active degrees of freedom because component entries outside their physical panels carry no mass.

Each full-order solver is verified before it is used to generate reduced-order data. The mechanical solver is compared with a finitely evaluated analytical Navier reference and an independent three-dimensional solid model. Its spatial behavior is then assessed through dual-reference convergence, against both the analytical plate solution and an independently refined unilateral full-order model, and through comparisons of two, three, four, and six-panel configurations with monolithic solutions. The thermomechanical solver is assessed through an auxiliary Navier-series comparison, plate-only refinement, coupled refinement, partitioned iteration histories, and a measured thermal and mechanical cost decomposition. The dynamic solver is checked through analytical and finite element modal comparisons, a spatially analytical single-mode transient reference, direct-to-modal consistency, undamped energy conservation, harmonic and localized-load response, and panel-to-monolithic agreement.

The limits of each evidence set are stated explicitly. Numerical agreement verifies the implemented approximation at the tested configurations and parameter values; it does not constitute experimental validation of a flooring product or calibration of panel connectors, support laws, heat-transfer coefficients, or damping [35, 36]. The finite element implementation uses the FEniCS/DOLFIN and Unified Form Language environments [37, 38].

## 1.4 Many-query reduction and reduced-order prediction

The governing problems are parameterized from the outset. Depending on the study, the physical parameters include orthotropic rigidities, foundation properties, load amplitudes or locations, environmental temperatures, solar flux, and heat-transfer coefficients. Repeated high-fidelity evaluation over these domains is expensive: the static problem requires nonlinear solves, the coupled thermomechanical problem requires repeated three-dimensional thermal and plate solves, and the transient problem requires a complete time integration for every physical query.

Proper Orthogonal Decomposition (POD) supplies the common linear-compression framework, and the resulting reduced-order models (ROMs) compute retained coordinates according to the state or output map of each study [39–43]. The mechanical study compares a non-hyper-reduced intrusive POD–Galerkin model with radial-basis-function interpolation of POD coordinates (PODI–RBF), linear interpolation of POD coordinates (PODI–Linear), Gaussian-process regression of POD coordinates (POD–GPR), and neural-network regression of POD coordinates (POD–NN). The intrusive method retains full-order residual and Jacobian assembly; the non-intrusive methods replace the online nonlinear solve by interpolation or regression.

Reduced-order treatments of parameterized three-dimensional heat transfer and steady one-way thermomechanical coupling provide relevant methodological context [44, 45]. The thermomechanical full-order model remains physically coupled, while its ROM assessment is field-specific. The same accepted converged coupled states generate two independent fields in the fixed output spaces used for POD and error evaluation: $\widehat{w}_h(\boldsymbol{\mu})$ and $\widehat{\theta}_h(\boldsymbol{\mu})$. Here, $\boldsymbol{\mu}$ denotes the case-specific physical-parameter vector, the subscript $h$ identifies the finite-dimensional full-order field, and the hats identify the corresponding output-space fields. The displacement and thermal-driver POD bases are constructed separately, with their retained dimensions selected and assessed in Chapter 8. The fields have different output spaces, physical units, POD metrics, bases, predictive models, and error measures while sharing the same case-specific sample assignments. The method inventory comprises a hybrid mechanical POD–Galerkin model, PODI–RBF, PODI–Linear, POD–GPR, POD–NN, a POD-enhanced deep-learning ROM (POD–DL-ROM), and field-specific direct deep-learning ROMs (DL-ROMs) [46–50]. The hybrid model reduces only the mechanical correction while retaining the three-dimensional thermal solve, thermal-to-plate reduction, and full-order mechanical assembly; it is therefore a partially reduced coupled workflow.

The transient reduced output is displacement only. Complete physical trajectories are assigned to disjoint training, validation, and test subsets before the stored states are flattened. Only training trajectories enter POD construction, preprocessing, and model fitting, so temporally adjacent states from one trajectory cannot cross statistical subsets. The retained dimensions are selected within the transient ROM assessment in Chapter 11. The active predictive methods are PODI–RBF, PODI–Linear, POD–GPR, POD–NN, and POD–DL-ROM. They map physical parameters and query time to displacement POD coordinates. Time $\tau$ is the independent evolution variable and a query coordinate of the surrogate, not a material or loading parameter. The resulting ROM is a time-conditioned, non-causal field predictor rather than a reduced Newmark propagator. Velocity, acceleration, stress resultants, reactions, bending moments, and interface jumps remain full-order quantities rather than active reduced outputs.

This output-aware viewpoint is especially important for panelized systems. Several ROM databases transfer a broken panel-wise displacement to a fixed continuous space before compression and field-error evaluation. The transfer preserves the influence of the seam mechanics on the FOM state but does not retain seam jumps as separate reduced quantities. Direct POD projection (POD–Proj) is likewise a non-predictive compression diagnostic because it requires an already available full-order field; it is not a predictive online method. Timing quantities are interpreted according to their measured operations. In particular, the reported transient online ratio compares one complete FOM trajectory with one reconstructed

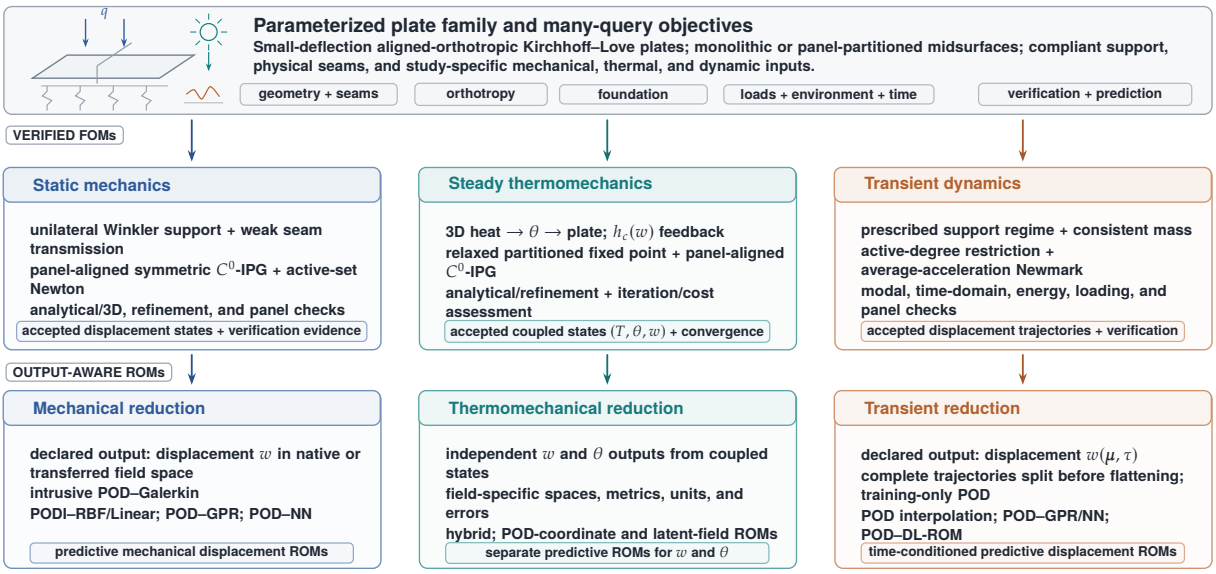



**Figure 1.2:** Unified model–FOM–ROM architecture of the dissertation. The mechanical, thermomechanical, and transient studies share the parameterized plate setting but retain separate governing equations and full-order solvers. Full-order states or trajectories are first mapped to the output used by the corresponding ROM. The mechanical branch contains intrusive and non-intrusive displacement ROMs; the thermomechanical branch uses separate displacement and thermal-driver one-field ROMs; and the transient branch uses complete-trajectory data splitting and time-conditioned displacement surrogates.

ROM state and is therefore a trajectory-to-single-state cost ratio rather than a matched trajectory-level speedup.

In Figure 1.2, $q$ denotes the transverse load, $T$ the three-dimensional temperature field, and $h_c(w)$ the displacement-dependent lower-face heat-transfer coefficient.

Figure 1.2 provides the thesis-level map of the three study chains; the detailed model, FOM, output, and ROM definitions are given in the corresponding study chapters.

## 1.5 Dissertation scope and contributions

The dissertation adopts small-deflection Kirchhoff–Love theory for plates with aligned orthotropy and local Winkler-type supports. It includes monolithic and panel-partitioned configurations, simply supported and free exterior boundaries, and mechanical, steady thermomechanical, and transverse transient loading. The static and thermomechanical plate problems use a regularized unilateral support law. In the transient study, one prescribed effective support coefficient is retained over each trajectory, yielding a linear evolution problem within the selected foundation regime. The thermomechanical and transient extensions are developed as separate physical studies, each with its own governing variables and output map. These choices define the scope of the full-order and reduced-order results.

Within this scope, the principal contributions are the following:

1. A unified parameterized plate framework is developed for three physical studies. They share aligned-orthotropic Kirchhoff–Love bending, compliant support interaction, monolithic or panel-partitioned midsurfaces, weak seam transmission, and a common upward-positive sign convention.

2. Panel-aligned symmetric $C^0$ interior-penalty FOMs are constructed while preserving the distinction between numerical stabilization inside a panel and mechanical coupling across a physical seam. The nonlinear static response is solved by an active-set Newton method with an algebraic residual criterion.
3. The mechanical implementation is supported by complementary reference evidence: Navier-series and independent three-dimensional comparisons, dual-reference $L^2(\Omega)$ and $H^1(\Omega)$ refinement, and two-, three-, four-, and six-panel assessments against monolithic solutions.
4. A physically coupled thermomechanical formulation and FOM connect steady three-dimensional heat conduction to plate bending through the thermal driver $\theta$. The partitioned solver, plate-only and coupled refinement, iteration histories, and thermal/mechanical computational-cost split are assessed without recasting the controlled tests as validation of the complete physical system.
5. A prescribed-regime transient model and verified FOM combine panel-aligned $C^0$-IPG stiffness, consistent transverse mass, active dynamic degrees of freedom, modal analysis, and average-acceleration Newmark integration. The verification program combines modal, time-domain, energy, loading, and panel-consistency checks within an explicitly fixed-discretization evidence scope.
6. Output-specific ROMs are developed rather than imposing one reduction on all three studies. The mechanical study compares intrusive and non-intrusive displacement ROMs. The thermomechanical study uses separate displacement and thermal-driver outputs. The transient-dynamics study uses displacement-only output, complete-trajectory separation before flattening, training-only POD construction, and time-conditioned non-causal coordinate surrogates.

Across the studies, the evidence is organized so that numerical verification, projection diagnostics, predictive assessment, and timing are interpreted according to the data, outputs, and measured operations that support them. Appendix A documents the software environment, data lineage, statistical partitions, reduced-model configurations, timing definitions, and repository organization associated with the reported computations.

The study chapters reuse the notation and structural assumptions established in Chapter 2, while introducing their study-specific physics, outputs, and numerical procedures.

## 1.6 Organization of the thesis

Chapter 2 presents the shared theoretical background: Kirchhoff–Love plate kinematics, aligned orthotropic bending, compliant supports and panel interfaces, variational structure, fourth-order finite element approximation, plate-level thermomechanical coupling, transverse dynamics, parameterized state and output maps, and POD-based model reduction. The study chapters provide the corresponding energies, algorithms, data partitions, and numerical settings.

Part I addresses static mechanics. Chapter 3 formulates the continuous orthotropic plate problem on a regularized unilateral Winkler foundation, including panel partitions, seam transmission, weak equilibrium, and parameter dependence. Chapter 4 develops and verifies the panel-wise

symmetric $C^0$ interior-penalty discretization and active-set Newton solver through analytical and three-dimensional comparisons, dual-reference refinement, and panel-partition studies. Chapter 5 constructs native- and transferred-space POD representations, compares intrusive POD–Galerkin with PODI–RBF, PODI–Linear, POD–GPR, and POD–NN for the monolithic benchmark, and assesses the non-intrusive models on panel-partitioned domains.

Part II contains the thermomechanical and transient studies. Chapter 6 formulates three-dimensional steady heat transfer, the through-thickness thermal driver, thermal bending, displacement-dependent feedback, and the coupled specialization. Chapter 7 discretizes the thermal and plate subproblems on independent meshes, reconstructs $\theta_h$, implements relaxed partitioned iteration, and assesses the auxiliary Navier comparison, plate-only and coupled refinement, iteration convergence, and computational cost. Chapter 8 develops the hybrid mechanical POD–Galerkin approximation and separate displacement-only and thermal-driver-only non-intrusive workflows with field-specific metrics, bases, and held-out predictive assessment.

Chapter 9 formulates prescribed-regime linear transient response with transverse translational inertia and time-dependent loading. Chapter 10 constructs the active-degree-restricted mass–stiffness system, implements direct and modal average-acceleration Newmark integration, and presents fixed-discretization verification. Chapter 11 transfers displacement histories to a common output space and assesses PODI–RBF, PODI–Linear, POD–GPR, POD–NN, and POD–DL-ROM through trajectory-wise separation, training-only POD construction, held-out validation, post-selection test-trajectory diagnostics, and trajectory-to-single-state timing ratios.

The concluding chapter synthesizes the three model–FOM–ROM hierarchies, states the evidence limitations, and identifies future directions. Appendix A documents the software environment, solver and output-transfer workflows, statistical partitions, reduced-model architectures, hyperparameters, post-processing conventions, timing definitions, and repository organization. Section A.5 summarizes the implemented ROM maps, neural architectures, and principal hyperparameters.

The next chapter fixes the common notation and theoretical structure used throughout the three studies.

# Theoretical Background for Parameterized Plate Models

# 2


2.1 Role and scope of the background chapter . . . . 10
2.2 Thin-plate idealization and midsurface variables . . . 10
2.3 Kirchhoff–Love kinematics and bending measures . . 11
2.4 Aligned orthotropic bending and energetic structure 12
2.5 Compliant supports and panel interfaces . . . . . 12
2.6 Variational structure and boundary conditions . . . 13
2.7 Discretization of the fourth-order plate operator 14
2.8 Thermomechanical coupling at plate level . . . . . 15
2.9 Transient dynamics and modal structure . . . . . 16
2.10 Parameterized states, outputs, and model reduction 17
2.11 Summary and handoff to the study chapters . . . . . 18


## 2.1 Role and scope of the background chapter

This chapter introduces the theoretical concepts shared by the three physical studies and their model–FOM–ROM hierarchies. It summarizes the thin-plate idealization, Kirchhoff–Love kinematics, aligned orthotropic bending, compliant support and panel-interface models, variational structure, fourth-order discretization, thermomechanical coupling, transient dynamics, parameterized state and output maps, and POD-based reduction [12–14, 23, 39, 40].

Chapters 3–5, Chapters 6–8, and Chapters 9–11 specialize this common framework to the static mechanical, thermomechanical, and transient studies. Their full-order and reduced-order chapters provide the corresponding discrete formulations, algorithms, parameter studies, data partitions, and numerical assessments.

## 2.2 Thin-plate idealization and midsurface variables

Let the reference midsurface be a bounded Lipschitz domain $\Omega \subset \mathbb{R}^2$, let $t > 0$ denote the plate thickness, and let $z$ measure signed distance from the midsurface. The corresponding three-dimensional reference body is

$$\Omega^t = \Omega \times \left(-\frac{t}{2}, \frac{t}{2}\right) \subset \mathbb{R}^3. \tag{2.1}$$

The geometric reduction and the associated normal-fiber kinematics are summarized in Figure 2.1. The sketch identifies the quantities introduced here and anticipates the Kirchhoff–Love restriction stated in Section 2.3.

A plate theory replaces the complete three-dimensional elastic description by midsurface variables and stress resultants obtained by integration through the thickness. This reduction is appropriate when the thickness is small relative to the in-plane dimensions and the response of interest is dominated by flexure [12–14].

The primary structural field in this dissertation is the transverse midsurface displacement $w$. Under the thesis-wide convention, $w > 0$ denotes upward displacement and a downward transverse load has negative sign. The models are small-deflection, bending-dominated Kirchhoff–Love models. They do not introduce independent in-plane displacement fields, membrane resultants, transverse-shear rotations, or geometric nonlinearities. These restrictions belong to the governing models and are inherited by the full-order and reduced-order approximations; they are not consequences of model reduction.

For a monolithic plate, $w$ is one scalar field on $\Omega$. A panelized assembly is represented by a non-overlapping partition and a collection $w = (w_i)_{i=1}^{M}$, with $w_i$ defined on panel $\Omega_i$. Independent panel traces are retained at a physical seam until transmission terms are applied. The broken

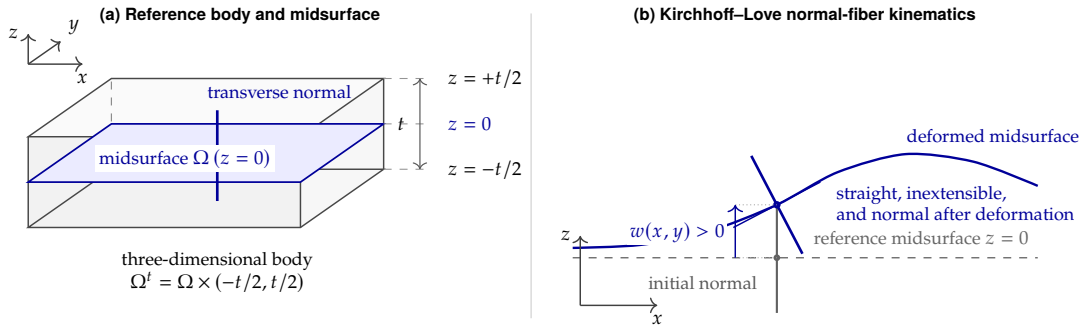



**Figure 2.1:** Thin-plate idealization used throughout the dissertation. Left: the three-dimensional reference body $\Omega^t = \Omega \times (-t/2, t/2)$, with $\Omega$ identified as the midsurface $z = 0$ and $z$ the signed thickness coordinate. Right: a representative normal section illustrating the Kirchhoff–Love restriction: a material fiber initially normal to the midsurface remains straight, inextensible, and normal after deformation, while the midsurface undergoes the upward-positive transverse displacement $w(x, y)$.

representation is therefore a model of a modular assembly rather than a mere subdivision of one globally continuous finite element field.

## 2.3 Kirchhoff–Love kinematics and bending measures

The Kirchhoff–Love hypothesis states that material lines initially normal to the midsurface remain straight, normal, and inextensible after deformation [12–14]. The right-hand panel of Figure 2.1 visualizes this restriction in a representative normal section. In the bending-only setting used here, the corresponding three-dimensional displacement field is represented by

$$\begin{aligned} u_x(x, y, z) &= -z\,\partial_x w(x, y), \qquad u_y(x, y, z) = -z\,\partial_y w(x, y), \\ u_z(x, y, z) &= w(x, y). \end{aligned} \tag{2.2}$$

Consequently, the transverse-shear strains vanish, rotations of a normal fiber are determined by first derivatives of $w$, and bending and twisting measures are determined by its second derivatives.

A convenient engineering-curvature vector is

$$\boldsymbol{\kappa}(w) := \begin{bmatrix} \partial_{xx} w \\ \partial_{yy} w \\ 2\partial_{xy} w \end{bmatrix}. \tag{2.3}$$

The factor two in the twisting component is an engineering convention that places the torsional rigidity in the third diagonal entry of the rigidity matrix. Some references adopt the opposite overall sign for curvature. Either convention is admissible provided that curvature, moment resultants, thermal moments, loading, and boundary resultants are defined consistently. The study chapters fix the signs used throughout the dissertation.

Because $\boldsymbol{\kappa}(w)$ contains second derivatives, the continuous bending problem naturally requires $w \in H^2(\Omega)$, or panel-wise $H^2$-regularity on a partition. This regularity requirement is the source of the principal finite element difficulty addressed by the $C^0$ interior-penalty formulations used later [23, 24].

## 2.4 Aligned orthotropic bending and energetic structure

Engineered panels can have different bending stiffnesses along two principal in-plane directions. The dissertation uses an aligned structurally orthotropic law: the orthotropy axes coincide with the Cartesian $x$- and $y$-directions [13, 14]. In compact notation, the curvature-conjugate moment vector is

$$\mathbf{m}(w) = \mathbf{D}\,\boldsymbol{\kappa}(w), \qquad \mathbf{D} = \begin{bmatrix} D_x & D_{xy} & 0 \\ D_{xy} & D_y & 0 \\ 0 & 0 & D_s \end{bmatrix}. \tag{2.4}$$

Here $D_x$ and $D_y$ are the principal flexural rigidities, $D_{xy}$ couples the two normal curvatures, and $D_s$ is the torsional rigidity. The constitutive coefficient $D_{xy}$ belongs to the orthotropic material law and is unrelated to the mechanical coupling between neighboring panels. The exact component definitions of $M_x$, $M_y$, and $M_{xy}$ are stated in Chapters 3, 6, and 9.

The elastic bending energy on a panel $A \subseteq \Omega$ is generated by the quadratic form

$$\mathcal{U}_{\mathrm{b},A}(w) = \frac{1}{2}\int_A \boldsymbol{\kappa}(w)^T \mathbf{D}\boldsymbol{\kappa}(w)\,\mathrm{d}\Omega. \tag{2.5}$$

A sufficient condition for positive definiteness is

$$D_x > 0, \qquad D_y > 0, \qquad D_s > 0, \qquad D_x D_y - D_{xy}^2 > 0. \tag{2.6}$$

The same elastic bending operator is reused by the static mechanical and transient-dynamics studies. In the thermomechanical study, thermally generated moments modify the total moment resultant but do not replace the elastic rigidity. This separation between constitutive stiffness and thermal loading is important both physically and computationally.

## 2.5 Compliant supports and panel interfaces

A Winkler foundation models the substrate as a continuous distribution of local springs: the reaction at one point depends on the displacement at that point rather than on a nonlocal substrate deformation field [15]. A bilateral linear law has the form $r_s(w) = k_s w$. In the static mechanical and thermomechanical studies, the active stiffness distinguishes compression from uplift, following the compression-only foundation setting [16]. A small positive uplift-side stiffness is introduced in this dissertation to regularize complete loss of contact. The support is linear on a fixed compression–uplift partition but nonlinear overall because the active partition depends on $w$. In the transient-dynamics study, one effective support coefficient is prescribed before time integration and remains fixed throughout the complete trajectory. That model is linear within the selected regime and does not represent transient contact, detachment, impact, recontact, or state switching.

Let the panelized midsurface satisfy

$$\overline{\Omega} = \bigcup_{i=1}^{M} \overline{\Omega_i}, \qquad \Omega_i \cap \Omega_j = \varnothing \quad (i \neq j). \tag{2.7}$$

If adjacent panels meet on $\Gamma_{ij} = \partial\Omega_i \cap \partial\Omega_j$, ideal Kirchhoff–Love kinematic transmission would require

$$w_i = w_j, \qquad \partial_n w_i = \partial_n w_j \quad \text{on } \Gamma_{ij}, \tag{2.8}$$

with a consistently oriented interface normal $n$. The first condition transmits transverse displacement; the second transmits normal rotation. The study models retain independent panel fields and enforce these conditions weakly through positive seam penalties. The finite penalties define the adopted panel model; they are not identified with experimentally calibrated connector laws.

Two notions of interface must remain separate. Element facets inside a physical panel are *numerical facets* introduced by the mesh. Their role is to stabilize a $C^0$ approximation of a fourth-order operator. Panel seams are *physical interfaces* introduced by the structural layout. Their role is to transmit kinematics between modular components. Related plate and shell couplings include penalty, mortar, Nitsche-type, and super-penalty constructions [28–32]. The exact mesh- and degree-dependent terms used in this dissertation are deferred to the full-order chapters.

## 2.6 Variational structure and boundary conditions

A variational formulation expresses equilibrium through virtual work or the stationarity of an energy functional. It identifies the admissible field space, incorporates essential boundary conditions into that space, and leaves the corresponding natural boundary conditions in the weak statement. For the panelized plate, the natural continuous space is broken and panel-wise $H^2$-regular. Let $V$ denote the corresponding monolithic or broken admissible $H^2$-type space, including the essential boundary conditions.

Let $a_{\boldsymbol{\mu}}$ denote elastic bending, $c_{\Gamma,\boldsymbol{\mu}}$ physical seam transmission, $g_{\boldsymbol{\mu}}$ the support contribution, and $\ell_{\boldsymbol{\mu}}$ external transverse loading. The static mechanical problem has the abstract form

$$a_{\boldsymbol{\mu}}(w, v) + c_{\Gamma,\boldsymbol{\mu}}(w, v) + g_{\boldsymbol{\mu}}(w; v) = \ell_{\boldsymbol{\mu}}(v) \qquad \forall v \in V. \tag{2.9}$$

The semicolon indicates that the regularized unilateral support is generally nonlinear in its state argument and linear in the test function.

The thermomechanical study augments the static plate structure by a thermal state and a thermal-to-plate map. Let $T$ denote the absolute temperature field, $T_{\mathrm{ref}}$ a constant stress-free reference temperature, and $\theta$ the resulting plate thermal driver. At an abstract level,

$$\mathcal{R}_T(T; w, \boldsymbol{\mu}) = 0, \qquad \theta = \mathcal{M}(T - T_{\mathrm{ref}}), \qquad \mathcal{R}_M(w; \theta, \boldsymbol{\mu}) = 0. \tag{2.10}$$

Here $\mathcal{R}_T$ is the thermal residual, $\mathcal{M}$ is the thermal-to-plate reduction map, and $\mathcal{R}_M$ is the mechanical residual including the thermal-bending functional. Dependence of the lower-face thermal exchange on $w$ closes a coupled fixed-point problem.

For a prescribed dynamic support regime, let $\tau$ denote time and let overdots denote differentiation with respect to $\tau$. The transient weak

problem has the form

$$m(\ddot{w}, v) + k_\mu(w, v) = \ell_\mu(\tau; v) \qquad \forall v \in V, \tag{2.11}$$

where $m$ is the transverse mass form and $k_\mu$ collects elastic bending, the fixed foundation coefficient, and seam transmission. The continuum model is undamped; optional Rayleigh damping is introduced only in the semidiscrete algebraic system, and the ROM campaigns use zero damping.

The exterior boundary decomposition used later distinguishes simply supported and free portions. On a simply supported edge, transverse displacement is an essential condition and the normal bending-moment condition is natural. On a free edge, the normal bending moment and effective Kirchhoff shear are natural conditions [12–14]. On polygonal domains, corner contributions are understood through the global variational statement rather than as independent pointwise edge equations.

For static elasticity, the weak form follows from stationarity of the total potential energy on each active branch. For dynamics, Hamilton's principle adds kinetic energy and yields the weak evolution equation. These principles explain the symmetry of the elastic and mass operators and provide a natural basis for finite element and projection-based approximations.

## 2.7 Discretization of the fourth-order plate operator

A conforming finite element approximation of the Kirchhoff–Love weak problem would require a discrete subspace of $H^2$, hence global $C^1$ continuity across element boundaries. Classical $C^1$ plate elements exist, but they are more restrictive than standard Lagrange elements on practical two-dimensional meshes [22–24]. The dissertation uses the symmetric $C^0$ interior-penalty approach [25–27].

For polynomial degree $p \geq 2$, the finite element displacement is continuous within each physical panel but its normal derivative may jump across element facets. The discrete plate form can be described schematically as

$$a_h = a_{\mathrm{pw}} + s_h^{\mathrm{cons}} + s_h^{\mathrm{sym}} + s_h^{\mathrm{pen}}, \tag{2.12}$$

where $a_{\mathrm{pw}}$ is the elementwise bending form, the next two terms restore consistency and symmetry, and the penalty controls jumps of the normal slope. These terms are assembled only on numerical facets inside a physical panel. The physical seam terms are added separately on $\Gamma_{ij}$. Exact moment contractions, trace conventions, and penalty scaling remain in the full-order chapters.

The resulting finite-dimensional problem is the high-fidelity or full-order model (FOM). In the mechanical study, its nonlinear residual is solved with an active-set Newton method. In the thermomechanical study, independent three-dimensional thermal and two-dimensional plate meshes are connected by numerical thickness integration and transfer to a plate thermal-driver space. In the transient-dynamics study, the plate stiffness is combined with a consistent transverse mass matrix. The global mixed realization of a panelized dynamic system contains component degrees of freedom outside their associated physical panels and hence with zero assembled mass contribution. A single active-index

restriction retains the degrees of freedom whose mass-diagonal entries exceed the numerical tolerance defined in Section 10.3; the same restriction is then used for eigenanalysis, time integration, lifting, and snapshot storage.

Numerical verification asks whether the discrete equations and algorithms solve the stated model at the tested settings. Experimental validation asks whether the model adequately represents a physical system [35, 36]. Analytical references, refined numerical references, fixed-discretization checks, and convergence studies support different claims and must be named accordingly. This distinction carries into model reduction: a reduced-order model (ROM) can approximate only the state or output map supplied by its parent FOM.

## 2.8 Thermomechanical coupling at plate level

The thermal extension begins from a steady heat-conduction problem on $\Omega^t$,

$$-\nabla \cdot (\mathbf{K}_T \nabla T) = s \qquad \text{in } \Omega^t, \tag{2.13}$$

where $T$ is the absolute temperature field, $s$ is a volumetric heat source, and $\mathbf{K}_T$ is the symmetric positive-definite thermal conductivity tensor. Their SI units are $[T] = \mathrm{K}$, $[s] = \mathrm{W\,m^{-3}}$, and $[\mathbf{K}_T] = \mathrm{W\,m^{-1}K^{-1}}$. The thermomechanical model combines convection and long-wave radiation on the exposed face, absorbed solar flux, displacement-dependent contact–gap exchange with the substrate on the lower face, and adiabatic lateral boundaries [17–20]. These mechanisms are introduced here only to identify the coupling path; their signs and nonlinear Robin terms are defined in Chapter 6.

A uniform temperature increment through the thickness produces no centered bending moment in the bending-only plate reduction. The relevant quantity is the first centered through-thickness moment of the temperature increment,

$$\theta(x,y) = \frac{12}{t^3} \int_{-t/2}^{t/2} z\big(T(x,y,z) - T_{\mathrm{ref}}\big)\,\mathrm{d}z, \qquad [\theta] = \mathrm{K\,m^{-1}}, \tag{2.14}$$

which defines the two-dimensional *thermal driver*. It is not a temperature. The normalization makes $\theta$ equal to the coefficient of a temperature field that varies linearly through the thickness. In the aligned orthotropic setting, the thermal bending vector has the abstract form

$$\mathbf{m}^{\mathrm{th}} = \mathbf{D}\,\mathbf{B}^{\mathrm{th}}\theta, \tag{2.15}$$

where $\mathbf{B}^{\mathrm{th}}$ contains the in-plane thermal-expansion coefficients. Chapter 6 fixes the sign with which this vector enters the total moment resultant.

If lower-face exchange is prescribed independently of the mechanical state, the heat and plate solves form a one-way sequence. When the exchange coefficient depends on the current displacement, the coupled iteration has the conceptual structure

$$w^{(k)} \longmapsto T^{(k+1)} \longmapsto \theta^{(k+1)} \longmapsto \widetilde{w}^{(k+1)} \longmapsto w^{(k+1)}. \tag{2.16}$$

The final step may include displacement relaxation. Radiation is resolved inside the thermal solve, the unilateral support active set inside the mechanical solve, and displacement–thermal feedback by the fixed-point

iteration. This is a partitioned coupled method rather than a monolithic thermomechanical Newton method [21]. A one-way calculation is the single-pass specialization of the same sequence.

The accepted coupled FOM state contains $(T_h, \theta_h, w_h)$. Chapter 8 treats two separate one-field outputs obtained from those same accepted states: the global displacement $\widehat{w}_h$ and the global thermal driver $\widehat{\theta}_h$. Their spaces, metrics, bases, predictors, units, and error measures remain separate, while the complete three-dimensional temperature field remains part of the full-order coupled workflow. The retained dimensions are selected and assessed in Chapter 8.

## 2.9 Transient dynamics and modal structure

The transient-dynamics study adds transverse inertia, initial data, and time-dependent loading to the elastic plate and seam structure. Time is denoted by $\tau$, while $t$ remains the plate thickness. For $u, v \in V$, the retained mass form is

$$m(u, v) = \int_\Omega \rho t \, uv \, \mathrm{d}\Omega, \tag{2.17}$$

where $\rho > 0$ is the volumetric mass density, $[\rho] = \mathrm{kg\,m^{-3}}$. Thus, the continuum model includes translational transverse inertia only. Rotary inertia associated with the in-plane components of Kirchhoff–Love motion is omitted, as are transverse-shear deformation and in-plane inertia. The initial-value problem requires $w(\cdot, 0) = w_0$ and $\dot{w}(\cdot, 0) = v_0$.

For each prescribed effective foundation coefficient, the stiffness is linear and time independent. In the absence of forcing and damping, harmonic separation leads to the generalized eigenproblem

$$k_{\boldsymbol{\mu}}(\varphi_n, v) = \omega_n^2 m(\varphi_n, v) \qquad \forall v \in V, \tag{2.18}$$

where $\omega_n$ is the $n$-th natural circular frequency and $\varphi_n$ is its mode. After spatial discretization,

$$\mathbf{K}_h \boldsymbol{\varphi}_{n,h} = \omega_n^2 \mathbf{M}_h \boldsymbol{\varphi}_{n,h}, \tag{2.19}$$

where $\mathbf{K}_h$ and $\mathbf{M}_h$ are the active finite element stiffness and mass matrices and $\boldsymbol{\varphi}_{n,h}$ is the discrete mode coefficient vector. The same modes support analytical frequency comparisons and modal superposition, while the reference trajectories are generated from the assembled finite element system [51].

The semidiscrete equation may be written as

$$\mathbf{M}_h \ddot{\mathbf{w}}_h(\tau) + \mathbf{C}_h \dot{\mathbf{w}}_h(\tau) + \mathbf{K}_h \mathbf{w}_h(\tau) = \mathbf{f}_h(\tau). \tag{2.20}$$

Here $\mathbf{w}_h$ and $\mathbf{f}_h$ are the active displacement and load coefficient vectors, and $\mathbf{C}_h$ is the optional algebraic Rayleigh-damping matrix [52]. In the ROM campaigns, $\mathbf{C}_h = \mathbf{0}$. The average-acceleration Newmark method advances the second-order system [34]. As stated in Section 2.5, the effective support coefficient is fixed during time integration.

A transient state depends on a physical trajectory parameter and observation time. Time remains the independent evolution variable rather than

a material, geometric, or loading parameter. A non-intrusive surrogate can nevertheless evaluate the map

$$(\boldsymbol{\mu}, \tau) \longmapsto \widehat{w}_h(\boldsymbol{\mu}, \tau). \tag{2.21}$$

The reported transient ROM is therefore time conditioned: it predicts the requested displacement field directly and does not define a causal transition $\mathbf{a}_n \mapsto \mathbf{a}_{n+1}$ or a reduced Newmark integrator. Displacement is the active reduced output; velocity, acceleration, moments, stresses, reactions, and seam jumps remain full-order quantities.

## 2.10 Parameterized states, outputs, and model reduction

Let $\mathcal{D} \subset \mathbb{R}^P$ be an admissible parameter domain. A parameter vector may contain material coefficients, support data, loading data, environmental inputs, or a selected subset of them. The FOM defines a high-dimensional state map, while the quantity reduced or observed may be only one output of that state. It is therefore useful to distinguish

$$\mathcal{S}_h : \mathcal{D} \longrightarrow \mathcal{X}_h, \qquad \mathcal{O}_h^{\chi} : \mathcal{X}_h \longrightarrow X_h^{\chi}, \qquad \widehat{\chi}_h(\boldsymbol{\mu}) = \mathcal{O}_h^{\chi}\big(\mathcal{S}_h(\boldsymbol{\mu})\big). \tag{2.22}$$

Here $\mathcal{X}_h$ is the complete FOM state space, $X_h^{\chi}$ is the space of the selected output $\chi$, and $\mathcal{O}_h^{\chi}$ may include field extraction or transfer. In the mechanical study, the state and displacement output may coincide. In the thermomechanical study, $\mathcal{S}_h(\boldsymbol{\mu}) = (T_h, \theta_h, w_h)$, while the active outputs are the two independent fields $\widehat{w}_h$ and $\widehat{\theta}_h$. For a panelized model, the output map can transfer a broken physical field to a fixed continuous space. This transfer provides a common coefficient ordering for compression, learning, and field-error evaluation. A separate output definition would be required to retain seam jumps independently. For transient dynamics, the domain of the state and output maps is $\mathcal{D} \times [0, \tau_{\mathrm{f}}]$, where $\tau_{\mathrm{f}} > 0$ is the final observation time; within this cross-study chapter, $T$ remains reserved for temperature.

For one output field, let $N_h^{\chi} := \dim(X_h^{\chi})$, let $N_s$ be the number of accepted snapshots, and let $\widehat{\boldsymbol{\chi}}_h^{(i)} \in \mathbb{R}^{N_h^{\chi}}$ denote the coefficient vector of snapshot $i$ in the selected output basis. The snapshot matrix is

$$\mathbf{S}_{\chi} = \left[\widehat{\boldsymbol{\chi}}_h^{(1)} \mid \cdots \mid \widehat{\boldsymbol{\chi}}_h^{(N_s)}\right] \in \mathbb{R}^{N_h^{\chi} \times N_s}. \tag{2.23}$$

Proper Orthogonal Decomposition (POD) identifies a low-dimensional linear space that minimizes aggregate projection error in a declared metric [40, 41, 53]. Let $N_{\chi}$ be the retained dimension, let $\mathbf{W}_{\chi}$ be the Gram matrix of the output metric, and let $\mathbf{V}_{N_{\chi}}$ contain the metric-orthonormal modes. With $\boldsymbol{\xi} = \boldsymbol{\mu}$ for static and steady outputs and $\boldsymbol{\xi} = (\boldsymbol{\mu}, \tau)$ for the transient output, the coefficient vector of an output field is approximated by

$$\widehat{\boldsymbol{\chi}}_h(\boldsymbol{\xi}) \approx \mathbf{V}_{N_{\chi}} \mathbf{z}_{N_{\chi}}(\boldsymbol{\xi}), \qquad \mathbf{V}_{N_{\chi}}^T \mathbf{W}_{\chi} \mathbf{V}_{N_{\chi}} = \mathbf{I}_{N_{\chi}}, \tag{2.24}$$

where $\mathbf{z}_{N_{\chi}}(\boldsymbol{\xi}) \in \mathbb{R}^{N_{\chi}}$ is the reduced-coordinate vector. The metric is part of the reduced output definition. The thermomechanical displacement field uses an $H^1$-type displacement-plus-gradient metric, whereas the thermal driver uses an $L^2(\Omega)$ metric. The transient-dynamics study first extracts an uncentered Euclidean singular subspace from the training snapshot matrix and then reorthonormalizes the spatial modes in an

$H^1$-type finite element product; its construction is therefore not described as a purely $H^1$-optimal POD.

The statistical partition precedes every data-dependent operation that could transfer information from held-out data into the model [54]. Training data define the basis, preprocessing, and predictive maps; validation data guide model selection; and test data are reserved for independent assessment. In the thermomechanical study, the displacement and thermal-driver workflows use the same case assignments while constructing separate training-based reduced spaces. For transient data, complete trajectories are assigned to statistical subsets before the stored states are flattened, so that states from one trajectory cannot cross subsets [55]. The campaign sizes, stored-time grids, and retained dimensions are reported in the corresponding reduced-order chapters.

Reduced-basis construction and surrogate response-surface approximation provide complementary computational-reduction strategies for parametrized PDEs [56].

The dissertation uses three related online strategies:

- In an *intrusive projection ROM*, the parent residual and, when required, its Jacobian are projected onto the reduced space. The mechanical POD–Galerkin model has this form. Without hyper-reduction, full-order assembly can remain in the online stage [57–59].
- In a *hybrid ROM*, only part of a coupled workflow is reduced. The thermomechanical hybrid model projects the mechanical correction while retaining the three-dimensional heat solve, thermal-to-plate reduction, and full-order mechanical assembly.
- In a *non-intrusive ROM*, interpolation or regression approximates the coordinate map $\boldsymbol{\xi} \mapsto \mathbf{z}_N(\boldsymbol{\xi})$ [46, 47, 60]. A POD-enhanced deep-learning ROM retains an initial POD compression and learns a nonlinear latent representation of the POD coordinates [49]. A direct deep-learning ROM instead learns and decodes a nonlinear latent representation of the output field without using POD coordinates as its predicted output [50, 61].

Projection of an available full-order field measures the compression error of the retained basis, whereas predictive assessment also includes the error of the learned parameter-to-coordinate or parameter-to-latent map. The study chapters therefore report projection error, predictive error, single-query timing, and complete-trajectory timing according to the operation being measured [39–41, 43]. The implemented ROM maps, neural architectures, and principal hyperparameters are summarized in Section A.5.

## 2.11 Summary and handoff to the study chapters

The common structural backbone of the dissertation is a small-deflection, aligned-orthotropic Kirchhoff–Love plate described by the transverse midsurface displacement. Second derivatives determine bending and twisting, a positive orthotropic rigidity matrix defines the elastic energy, local Winkler-type reactions represent compliant support, and broken panel fields permit weak transmission of displacement and normal rotation across physical seams.

The three physical studies specialize this backbone in different directions. The static mechanical and thermomechanical models use a regularized unilateral support law. The thermomechanical model maps a steady three-dimensional temperature field to the plate thermal driver and may close a displacement-dependent fixed point; its ROM assessment uses separate displacement and thermal-driver representations. The transient model adds translational inertia and time-dependent loading while fixing the effective support regime over each trajectory. Its reported surrogate predicts displacement at a requested parameter–time pair.

The fourth-order spatial operator motivates panel-aligned symmetric $C^0$ interior-penalty discretization, with numerical element facets kept distinct from physical seams. The resulting FOM states are mapped to declared outputs before compression and prediction. POD defines reduced representations; intrusive projection retains the parent operators, hybrid reduction compresses only selected subproblems, and non-intrusive methods learn parameter-to-coordinate or parameter-to-latent maps.

Chapter 3 next fixes the exact static mechanical model. Chapters 3–11 form the complete study sequence, carrying the mechanical, thermomechanical, and transient formulations from continuous models to assessed or verified FOMs and output-specific ROMs. The corresponding residuals, penalty prefactors, algorithms, parameter domains, data counts, numerical results, and assessment settings are specified in the respective study chapters.

Part I

# Static Mechanics of Orthotropic Plates

> Beginning with a parameterized orthotropic Kirchhoff–Love plate on a regularized unilateral foundation, this part develops the continuous mechanics, verifies a panel-aligned $C^0$ interior-penalty full-order model against analytical, three-dimensional, refined, and monolithic references, and compresses the resulting displacement fields through intrusive and non-intrusive ROMs. It connects *orthotropic plate mechanics and weak seam transmission* with *many-query prediction*.

# Parameterized Mechanical Kirchhoff–Love Plate Model

# 3


3.1 Role of the mechanical model in the thesis . . . . 21
3.2 Geometric setting and panel partition . . . . . 21
3.3 Kirchhoff–Love kinematics and sign conventions . . . 22
3.4 Orthotropic bending resultants . . . . . 23
3.5 Unilateral Winkler foundation and uplift regularization . . . . . 23
3.6 Interface transmission across panel partitions . . 24
3.7 Energetic formulation . . 24
3.8 Weak formulation . . . . . 25
3.9 Parameter dependence and nonlinear structure . . . . 26
3.10 Summary of the mechanical model . . . . . 27


## 3.1 Role of the mechanical model in the thesis

This chapter develops the continuous mechanical plate model that underlies the high-fidelity approximation in Chapter 4 and the mechanical reduced-order models in Chapter 5. The focus here is the static mechanical response of thin panelized flooring systems.

The setting follows the mechanical viewpoint established in Chapter 1. The structure is idealized as a collection of thin, structurally orthotropic Kirchhoff–Love plates resting on a Winkler-type elastic foundation, using standard plate-theory terminology and assumptions [12–14]. The panel layout is represented by a non-overlapping partition of the plate midsurface, and the kinematic transmission across adjacent panels is enforced through penalty terms acting on displacement and normal rotation. The purpose of the chapter is to formulate this problem at the continuous level, identify the energetic contributions, state the associated weak problem, and make explicit the parameter dependence that enters the high-fidelity discretization in Chapter 4 and the reduced-order construction in Chapter 5.

## 3.2 Geometric setting and panel partition

Let $(\mathbf{e}_x, \mathbf{e}_y, \mathbf{e}_z)$ be an orthonormal basis of $\mathbb{R}^3$, and let $(x, y, z)$ denote the associated Cartesian coordinates. In the reference configuration, the plate midsurface occupies a bounded open set $\Omega \subset \mathbb{R}^2$ with Lipschitz boundary $\partial\Omega$. The coordinates $(x, y)$ identify material points on the midsurface, while the $z$-direction is transverse to the plate.

The geometric setting used throughout this chapter is shown schematically in Figure 3.1. The figure distinguishes the monolithic plate domain from a representative non-overlapping panel partition and introduces the interface notation used below.

A monolithic panel corresponds to $M = 1$. More general panelized layouts are represented by a finite collection of open, non-overlapping subdomains $\{\Omega_i\}_{i=1}^{M}$ such that

$$\overline{\Omega} = \bigcup_{i=1}^{M} \overline{\Omega_i}, \qquad \Omega_i \cap \Omega_j = \emptyset \quad \text{for } i \neq j.$$

For two adjacent panels, the common interface is denoted by

$$\Gamma_{ij} := \partial\Omega_i \cap \partial\Omega_j,$$

and the internal skeleton is

$$\Gamma := \bigcup_{1 \le i < j \le M} \Gamma_{ij}.$$

On each interface $\Gamma_{ij}$, the outward unit normals to $\Omega_i$ and $\Omega_j$ are denoted by $\mathbf{n}_i$ and $\mathbf{n}_j$, respectively. For defining interface jumps, a reference

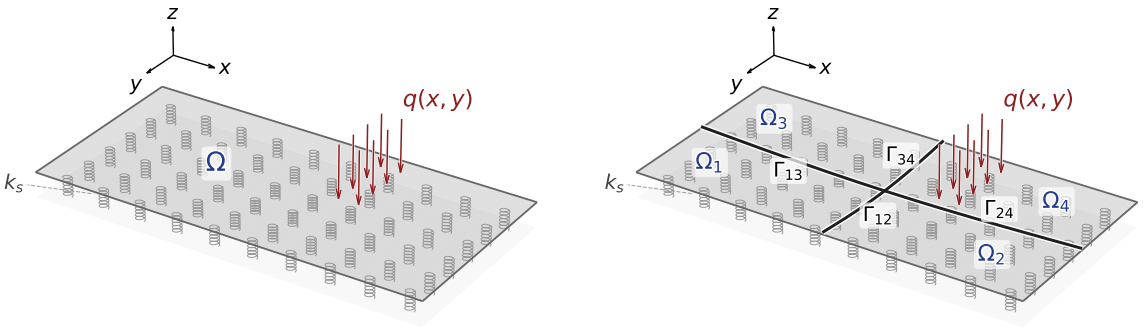



**Figure 3.1:** Geometric setting for the mechanical Kirchhoff–Love plate model on a Winkler-type elastic foundation. Left: monolithic plate domain $\Omega$ subjected to a distributed transverse load $q(x, y)$. Right: representative non-overlapping partition $\overline{\Omega} = \bigcup_{i=1}^{M} \overline{\Omega_i}$ with internal interfaces $\Gamma_{ij}$.

normal $\mathbf{n}$ is fixed on each $\Gamma_{ij}$, with

$$\mathbf{n} = \mathbf{n}_i = -\mathbf{n}_j.$$

The exterior boundary is decomposed as

$$\partial\Omega = \Gamma_S \cup \Gamma_F, \qquad \Gamma_S \cap \Gamma_F = \emptyset,$$

where $\Gamma_S$ denotes the simply supported portion and $\Gamma_F$ the free portion. This notation is retained in Chapter 4, where the essential part of the simply supported boundary condition is imposed directly in the discrete space.

## 3.3 Kirchhoff–Love kinematics and sign conventions

The unknown mechanical field is the transverse displacement of the midsurface. In the multi-panel setting, this field is represented as the broken scalar field

$$w = \{w_i\}_{i=1}^{M}, \qquad w_i : \Omega_i \to \mathbb{R}.$$

Throughout this chapter, $w > 0$ denotes upward transverse displacement. The transverse load $q(x, y)$ is taken positive in the same upward direction. Consequently, a downward load is represented by $q < 0$.

The Kirchhoff–Love assumption is used at the plate level: transverse normals to the midsurface remain straight and normal during deformation, and transverse shear deformation is neglected [62]. In the present static bending setting, this reduces the mechanical description to the scalar field $w$ on the midsurface. The bending response is governed by second derivatives of $w_i$ on each panel $\Omega_i$, so the continuous model naturally involves $H^2$-regularity panel-wise. The need to approximate this fourth-order structure with $C^0$ finite element spaces is one of the central motivations for the high-fidelity discretization in Chapter 4.

## 3.4 Orthotropic bending resultants

The orthotropic axes of each panel are assumed to coincide with the Cartesian axes $x$ and $y$. The bending response is described by the flexural rigidities $D_x$, $D_y$, and $D_{xy}$, together with the torsional rigidity $D_s$. These quantities may represent intrinsic material orthotropy or an effective orthotropy induced by the construction of the panel.

For a displacement field $w_i$ on $\Omega_i$, the bending and twisting moment resultants are defined by

$$\begin{aligned} M_x(w_i) &= D_x\,\partial_{xx}^2 w_i + D_{xy}\,\partial_{yy}^2 w_i, \\ M_y(w_i) &= D_y\,\partial_{yy}^2 w_i + D_{xy}\,\partial_{xx}^2 w_i, \\ M_{xy}(w_i) &= 2D_s\,\partial_{xy}^2 w_i. \end{aligned} \tag{3.1}$$

The corresponding bending energy stored in panel $\Omega_i$ is

$$\begin{aligned} \mathcal{U}_{\mathrm{p},i}(w_i;\boldsymbol{\mu}) &= \frac{1}{2}\int_{\Omega_i}\Big[M_x(w_i)\,\partial_{xx}^2 w_i + M_y(w_i)\,\partial_{yy}^2 w_i \\ &\qquad + 2M_{xy}(w_i)\,\partial_{xy}^2 w_i\Big]\,\mathrm{d}\Omega \\ &= \frac{1}{2}\int_{\Omega_i}\Big[D_x\big(\partial_{xx}^2 w_i\big)^2 + D_y\big(\partial_{yy}^2 w_i\big)^2 \\ &\qquad + 2D_{xy}\,\partial_{xx}^2 w_i\,\partial_{yy}^2 w_i + 4D_s\big(\partial_{xy}^2 w_i\big)^2\Big]\,\mathrm{d}\Omega. \end{aligned} \tag{3.2}$$

The rigidities are assumed to be chosen so that the bending energy is positive on admissible curvature fields. In particular, the computations in Chapter 4 use positive flexural and torsional rigidities and an admissible coupling rigidity.

## 3.5 Unilateral Winkler foundation and uplift regularization

The support is modeled as a Winkler-type foundation, so that the substrate response is local in the transverse displacement [63]. The present mechanical model uses a unilateral form of this law to distinguish compression from uplift. Under the upward-positive sign convention, compression of the support corresponds to $w_i < 0$, while uplift corresponds to $w_i \geq 0$.

For each panel $\Omega_i$, the foundation contribution to the total potential energy is

$$\Psi_{\mathrm{f},i}(w_i;\boldsymbol{\mu}) = \frac{1}{2}\int_{\Omega_i} k_s\left(\chi_{\{w_i<0\}} + \varepsilon\,\chi_{\{w_i\geq 0\}}\right) w_i^2\,\mathrm{d}\Omega, \tag{3.3}$$

where $k_s > 0$ is the foundation stiffness, $\chi$ denotes the indicator function, and $0 < \varepsilon \ll 1$ is a small uplift-side regularization parameter. Thus, the support acts with stiffness $k_s$ in compression and with the reduced stiffness $\varepsilon k_s$ in uplift. The ideal unilateral foundation law is recovered formally as $\varepsilon \to 0^+$ [64].

The unilateral foundation term is the source of nonlinearity in the static mechanical model. For a fixed active set, namely for a fixed partition of the plate into compression and uplift regions, the foundation contribution

is linear in $w_i$. The active set itself, however, depends on the unknown displacement and must therefore be handled in the nonlinear solution strategy of the high-fidelity approximation.

**Remark 3.5.1** (Active-set interpretation) The foundation contribution is linear for a fixed compression-uplift partition, but the partition itself depends on the unknown displacement. The nonlinearity of the mechanical problem therefore enters through the sign-dependent active set, not through the bending, interface, or loading terms.

## 3.6 Interface transmission across panel partitions

When adjacent panels meet along an internal interface, the mechanical model must transmit the relevant Kirchhoff–Love kinematic information across the seam. In this work, continuity of transverse displacement and continuity of normal rotation are imposed weakly through interface penalties [65].

For a broken scalar field $v = \{v_i\}_{i=1}^{M}$, the normal derivative on $\Gamma_{ij}$ is taken with respect to the fixed reference normal $\mathbf{n}$:

$$\partial_{\mathbf{n}} v_i := \mathbf{n} \cdot \nabla v_i .$$

The displacement and normal-rotation jumps across $\Gamma_{ij}$ are defined by

$$[\![v]\!]_{ij} := v_i - v_j, \qquad [\![\partial_{\mathbf{n}} v]\!]_{ij} := \partial_{\mathbf{n}} v_i - \partial_{\mathbf{n}} v_j .$$

The interface penalty energy is then

$$\Pi_{\mathrm{int}}(w;\boldsymbol{\mu}) = \sum_{1\le i<j\le M} \frac{1}{2} \int_{\Gamma_{ij}} \left[ \eta_w \, [\![w]\!]_{ij}^2 + \eta_\theta \, [\![\partial_{\mathbf{n}} w]\!]_{ij}^2 \right] \mathrm{d}\Gamma, \tag{3.4}$$

where $\eta_w > 0$ and $\eta_\theta > 0$ are interface weights associated with displacement and normal-rotation continuity, respectively. These parameters are part of the mechanical coupling model and are specified at the discrete level in Chapter 4 through mesh- and degree-dependent scaling.

## 3.7 Energetic formulation

The total potential energy combines the panel-wise bending energy, the unilateral foundation energy, the interface penalty energy, and the work of the transverse loading. For a prescribed load $q(x, y)$, the external work over panel $\Omega_i$ is

$$\mathcal{W}_{\mathrm{ext},i}(w_i;\boldsymbol{\mu}) = \int_{\Omega_i} q \, w_i \, \mathrm{d}\Omega. \tag{3.5}$$

The load $q$ may represent several mechanical loading configurations. The cases used later include a uniform load $q(x, y) = f$, patch loads of the form

$$q(x,y) = \sum_{k=1}^{N_{\mathrm{P}}} f_k \, \chi_{\mathcal{P}_k}(x,y),$$

where $\mathcal{P}_k \subset \Omega$ are loaded patches, and linearly varying loads of the form

$$q(x,y) = f_s + (f_e - f_s)\frac{x - x_s}{x_e - x_s}.$$

Here $f$, $f_k$, $f_s$, and $f_e$ follow the upward-positive sign convention.

The total potential energy is

$$\Pi(w;\boldsymbol{\mu}) = \sum_{i=1}^{M} \left[\mathcal{U}_{\mathrm{p},i}(w_i;\boldsymbol{\mu}) + \Psi_{\mathrm{f},i}(w_i;\boldsymbol{\mu}) - \mathcal{W}_{\mathrm{ext},i}(w_i;\boldsymbol{\mu})\right] + \Pi_{\mathrm{int}}(w;\boldsymbol{\mu}). \tag{3.6}$$

The model may depend on the rigidities, foundation stiffness, uplift regularization, interface weights, and load data. The parameter vector $\boldsymbol{\mu}$ denotes the subset varied in a particular study; the remaining quantities are fixed configuration or discretization data.

## 3.8 Weak formulation

To state the variational problem, define the broken admissible space

$$V := \left\{ v = \{v_i\}_{i=1}^{M} \;\middle|\; v_i \in H^2(\Omega_i),\ i = 1,\dots,M,\ v_i = 0 \text{ on } \Gamma_S \cap \partial\Omega_i \right\}.$$

The condition $v_i = 0$ on $\Gamma_S \cap \partial\Omega_i$ incorporates the essential part of the simply supported boundary condition. The bending-moment condition on $\Gamma_S$ and the free-edge conditions on $\Gamma_F$ arise naturally from the variational formulation.

The equilibrium displacement is characterized by stationarity of the total potential energy:

$$\delta_v \Pi(w;\boldsymbol{\mu}) = 0 \qquad \forall\, v \in V.$$

**Continuous mechanical weak problem**

Given $\boldsymbol{\mu} \in \mathcal{D}$, find $w \in V$ such that

$$a^{\boldsymbol{\mu}}(w,v) + g^{\boldsymbol{\mu}}(w;v) + c^{\boldsymbol{\mu}}(w,v) = \ell^{\boldsymbol{\mu}}(v) \qquad \forall\, v \in V. \tag{3.7}$$

The bending form is

$$a^{\boldsymbol{\mu}}(w,v) = \sum_{i=1}^{M} \int_{\Omega_i} \Big[ D_x\, \partial^2_{xx} w_i\, \partial^2_{xx} v_i + D_y\, \partial^2_{yy} w_i\, \partial^2_{yy} v_i + D_{xy} \left( \partial^2_{xx} w_i\, \partial^2_{yy} v_i + \partial^2_{yy} w_i\, \partial^2_{xx} v_i \right) + 4 D_s\, \partial^2_{xy} w_i\, \partial^2_{xy} v_i \Big] \, \mathrm{d}\Omega.$$

The interface form is

$$c^{\boldsymbol{\mu}}(w,v) = \sum_{1 \le i < j \le M} \int_{\Gamma_{ij}} \left[ \eta_w\, [\![ w ]\!]_{ij}\, [\![ v ]\!]_{ij} + \eta_\theta\, [\![ \partial_{\mathbf{n}} w ]\!]_{ij}\, [\![ \partial_{\mathbf{n}} v ]\!]_{ij} \right] \mathrm{d}\Gamma.$$

The foundation term is nonlinear in its first argument and linear in the test function:

$$g^{\boldsymbol{\mu}}(w;v) = \sum_{i=1}^{M} \int_{\Omega_i} k_s \left( \chi_{\{w_i < 0\}} + \varepsilon\, \chi_{\{w_i \ge 0\}} \right) w_i\, v_i \, \mathrm{d}\Omega.$$

Finally, the load functional is

$$\ell^{\mu}(v) = \sum_{i=1}^{M} \int_{\Omega_i} q\, v_i \, \mathrm{d}\Omega.$$

For completeness, the classical boundary resultants associated with smooth solutions are recorded here. Let $\mathbf{n} = (n_x, n_y)$ be the outward unit normal on $\partial\Omega$, let $\mathbf{t} = (-n_y, n_x)$ be the unit tangent, and define

$$\begin{aligned} M_{nn}(w) &:= M_x(w)n_x^2 + 2M_{xy}(w)n_x n_y + M_y(w)n_y^2, \\ M_{nt}(w) &:= \big(M_y(w) - M_x(w)\big)n_x n_y + M_{xy}(w)\big(n_x^2 - n_y^2\big). \end{aligned}$$

With

$$Q_x(w) := \partial_x M_x(w) + \partial_y M_{xy}(w), \qquad Q_y(w) := \partial_x M_{xy}(w) + \partial_y M_y(w),$$

the normal shear and effective Kirchhoff shear are

$$Q_n(w) := Q_x(w)n_x + Q_y(w)n_y, \qquad V_n(w) := Q_n(w) + \partial_s M_{nt}(w),$$

where $\partial_s$ denotes differentiation along $\mathbf{t}$. On smooth boundary portions, the exterior boundary conditions are

$$\begin{aligned} w = 0, \quad M_{nn}(w) = 0 \quad &\text{on } \Gamma_S, \\ M_{nn}(w) = 0, \quad V_n(w) = 0 \quad &\text{on } \Gamma_F. \end{aligned}$$

These conditions clarify the physical boundary interpretation of the weak formulation. The discrete formulation in Chapter 4 imposes $w_h = 0$ on $\Gamma_S$ strongly and retains the remaining boundary conditions in natural form.

## 3.9 Parameter dependence and nonlinear structure

The continuous mechanical model depends on the orthotropic rigidities, foundation stiffness, uplift regularization, interface weights, and load data entering the forms above. In the reported mechanical parameter studies, $\varepsilon$ is fixed and the interface weights are prescribed by the mesh- and degree-dependent rules of Section 4.4; they are not sampled coordinates.

A representative sampled parameter vector is

$$\boldsymbol{\mu} = (D_x, D_y, D_{xy}, D_s, k_s, \boldsymbol{\lambda}_q) \in \mathcal{D} \subset \mathbb{R}^P,$$

where $\boldsymbol{\lambda}_q$ denotes the finite-dimensional collection of load parameters. For a uniform load,

$$\boldsymbol{\lambda}_q = f, \qquad \boldsymbol{\mu} = (D_x, D_y, D_{xy}, D_s, k_s, f).$$

The numerical studies in Chapter 4 vary the orthotropic rigidities, foundation stiffness, and loading data. Separate verification studies vary the polynomial degree, mesh resolution, and panel layout as discretization or configuration choices rather than as coordinates of $\boldsymbol{\mu}$. The same sampled physical parameters form the input space of the reduced-order models in Chapter 5.

For fixed parameters, the bending, interface, and external-load contributions are linear. The unilateral foundation retains the active-set nonlinearity described in Section 3.5, so the discrete tangent operator depends on the current compression–uplift partition.

## 3.10 Summary of the mechanical model

This chapter has formulated the static mechanical model used in Part I of the thesis. The model combines orthotropic Kirchhoff–Love bending, unilateral Winkler-type support interaction, weak panel-interface transmission, and explicit parameter dependence. Its continuous statement is energetic and variational: the equilibrium displacement is the stationary point of the total potential energy (3.6), equivalently the solution of the weak problem (3.7).

Chapter 4 discretizes this model by a symmetric $C^0$ interior-penalty finite element method. The discretization preserves the panel-wise structure of the formulation, introduces stabilization terms for the fourth-order plate operator inside each panel, and treats the panel interfaces through penalty terms consistent with the continuous coupling energy. Chapter 5 then uses the resulting high-fidelity solutions to construct and assess reduced-order models for the same parameterized mechanical problem.

# Full-Order Approximation of the Mechanical Plate Model

# 4


4.1 Role of the high-fidelity approximation . . . . . 28
4.2 Discrete spaces and panel-wise finite element meshes 28
4.3 $C^0$ interior-penalty discretization inside panels . 29
4.4 Weak enforcement of panel-interface continuity 30
4.5 Discrete unilateral foundation term . . . . . 31
4.6 Mechanical full-order residual problem . . . . . 31
4.7 Nonlinear solution strategy . . . . . 32
4.8 Verification against analytical and three-dimensional references . . . . . 33
4.9 Mesh-refinement convergence study . . . . . 37
4.10 Penalty-coupled panel-partition assessment . . . 39
4.11 Summary of the full-order model . . . . . 41


## 4.1 Role of the high-fidelity approximation

Chapter 3 formulated the static mechanical model for parameterized orthotropic Kirchhoff–Love plates on unilateral Winkler-type foundations. The continuous problem is fourth order on each panel, nonlinear through the unilateral foundation response, and weakly coupled across panel interfaces through displacement and normal-rotation penalties. This chapter constructs the corresponding high-fidelity finite element approximation, based on the $C^0$ interior-penalty treatment of fourth-order plate problems [25, 27], and verifies it on the mechanical benchmark problems used to support the reduced-order constructions in Chapter 5.

In the terminology of Chapter 5, this high-fidelity discretization is the full-order model (FOM): it is the unreduced finite element problem from which reference solutions and snapshot data are generated.

The development proceeds from the $C^0$ interior-penalty Galerkin discretization of the fourth-order Kirchhoff–Love operator on standard $H^1$-conforming Lagrange spaces, through the discrete nonlinear residual problem and the Newton active-set strategy for the unilateral foundation law, to the mechanical verification studies. These compare the approximation against analytical and three-dimensional references for a monolithic simply supported plate, examine mesh-refinement convergence for isotropic-compatible and orthotropic rigidity configurations, and assess penalty-coupled panel partitions.

## 4.2 Discrete spaces and panel-wise finite element meshes

Let

$$\overline{\Omega} = \bigcup_{i=1}^{M} \overline{\Omega_i}$$

be the non-overlapping panel partition introduced in Section 3.2. On each panel $\Omega_i$, let $\mathcal{T}_h^i$ be a conforming, shape-regular triangulation with characteristic size $h_i$. For a polynomial degree $p \geq 2$, define

$$W_h^i := \left\{ \phi_h \in C^0(\overline{\Omega_i}) \;\middle|\; \phi_h|_T \in \mathbb{P}_p(T) \quad \forall\, T \in \mathcal{T}_h^i \right\}.$$

The global broken finite element space is the direct sum

$$W_h := \bigoplus_{i=1}^{M} W_h^i.$$

The direct-sum construction means that degrees of freedom belonging to different panels remain independent before the interface penalties are assembled. Thus, displacement continuity across physical panel interfaces is not imposed strongly.

The essential part of the simply supported boundary condition is imposed by the constrained space

$$V_h := \{w_h \in W_h \mid w_h = 0 \text{ on } \Gamma_S\}.$$

The remaining simply supported bending-moment condition and the free-edge conditions are retained in natural form.

## 4.3 $C^0$ interior-penalty discretization inside panels

Because $V_h$ is only $C^0$, it is not a conforming subspace of the panel-wise $H^2$ space required by the Kirchhoff–Love weak form. The missing regularity is restored weakly by a symmetric $C^0$ interior-penalty stabilization acting on mesh edges strictly inside each panel. This stabilization penalizes jumps of the normal derivative across element interfaces while preserving the use of standard Lagrange finite elements [25, 27, 66, 67].

For $\phi_h \in V_h$, define the discrete curvature vector and the orthotropic bending matrix by

$$\boldsymbol{\kappa}(\phi_h) := \begin{bmatrix} -\partial_{xx}\phi_h \\ -\partial_{yy}\phi_h \\ -2\partial_{xy}\phi_h \end{bmatrix}, \qquad \mathbf{D} := \begin{bmatrix} D_x & D_{xy} & 0 \\ D_{xy} & D_y & 0 \\ 0 & 0 & D_s \end{bmatrix}.$$

All second derivatives in this section are understood elementwise. The elementwise bending contribution over panel $\Omega_i$ is

$$b_i(w_h, \phi_h) := \int_{\Omega_i} \boldsymbol{\kappa}(w_h)^T \mathbf{D}\, \boldsymbol{\kappa}(\phi_h)\, \mathrm{d}\Omega. \tag{4.1}$$

Let $E = \partial T^+ \cap \partial T^-$ be a mesh edge shared by two neighboring elements $T^+, T^- \in \mathcal{T}_h^i$, and let $\boldsymbol{\nu}_E = (\nu_x, \nu_y)$ be a fixed unit normal on $E$, oriented from $T^-$ to $T^+$. The normal derivative and its jump are

$$\partial_{\boldsymbol{\nu}_E}\phi_h := \boldsymbol{\nu}_E \cdot \nabla\phi_h, \qquad \left[\!\left[\partial_{\boldsymbol{\nu}_E}\phi_h\right]\!\right] := \partial_{\boldsymbol{\nu}_E}(\phi_h|_{T^+}) - \partial_{\boldsymbol{\nu}_E}(\phi_h|_{T^-}).$$

Let

$$E_h^{i,\mathrm{int}} := \left\{E \text{ mesh edge of } \mathcal{T}_h^i \ : \ E \subset \Omega_i\right\}$$

be the set of edges strictly interior to panel $\Omega_i$. Since functions in $W_h^i$ are continuous inside the panel, the displacement jump across such edges is zero, while the normal derivative may be discontinuous.

For $E \in E_h^{i,\mathrm{int}}$, define the normal-normal bending moment

$$m_{\boldsymbol{\nu}_E}(\phi_h) := M_x(\phi_h)\nu_x^2 + 2M_{xy}(\phi_h)\nu_x\nu_y + M_y(\phi_h)\nu_y^2,$$

where the moment resultants are those in (3.1). Its average across $E$ is

$$\{\{m_{\boldsymbol{\nu}_E}(\phi_h)\}\} := \frac{1}{2}\left(m_{\boldsymbol{\nu}_E}(\phi_h|_{T^+}) + m_{\boldsymbol{\nu}_E}(\phi_h|_{T^-})\right).$$

With the effective stiffness scale

$$D_{\mathrm{eff}} := \sqrt{D_x D_y}, \qquad \tau_E := \frac{(p+2)^2 D_{\mathrm{eff}}}{h_E},$$

where $h_E$ is the facetwise average of the diameters of the two elements adjacent to $E$ [68], the panel-wise $C^0$ interior-penalty stabilization is

$$\begin{aligned} j_i(w_h, \phi_h) := \sum_{E \in E_h^{i,\mathrm{int}}} \Bigg[ & - \int_E \{\{m_{\boldsymbol{\nu}_E}(w_h)\}\} \left[\left[\partial_{\boldsymbol{\nu}_E} \phi_h\right]\right] \, \mathrm{d}s \\ & - \int_E \{\{m_{\boldsymbol{\nu}_E}(\phi_h)\}\} \left[\left[\partial_{\boldsymbol{\nu}_E} w_h\right]\right] \, \mathrm{d}s \\ & + \int_E \tau_E \left[\left[\partial_{\boldsymbol{\nu}_E} w_h\right]\right] \left[\left[\partial_{\boldsymbol{\nu}_E} \phi_h\right]\right] \, \mathrm{d}s \Bigg]. \end{aligned} \tag{4.2}$$

The global bending and panel-interior stabilization forms are

$$a_h^{\mu}(w_h, \phi_h) := \sum_{i=1}^{M} b_i(w_h, \phi_h), \qquad s_h^{\mu}(w_h, \phi_h) := \sum_{i=1}^{M} j_i(w_h, \phi_h).$$

On the exterior boundary, the displacement condition on $\Gamma_S$ is already imposed strongly through $V_h$, while the associated normal-moment condition remains natural. No exterior penalty is added on $\Gamma_F$; the free-edge moment and effective-shear conditions are likewise retained in natural form. The facet terms in (4.2) therefore act only on numerical mesh facets strictly inside a physical panel.

## 4.4 Weak enforcement of panel-interface continuity

The panel interfaces are treated separately from the numerical mesh facets interior to each panel. Across a physical seam $\Gamma_{ij}$, the discrete coupling is the direct counterpart of the interface energy in Equation (3.4). Written first in terms of positive facetwise weights $\eta_w$ and $\eta_\theta$, the interface form is

$$\begin{aligned} c_{\Gamma,h}^{\mu}(w_h, \phi_h) := \sum_{1 \le i < j \le M} \int_{\Gamma_{ij}} \Big[ & \eta_w \left[\!\left[ w_h \right]\!\right]_{ij} \left[\left[\phi_h\right]\right]_{ij} \\ & + \eta_\theta \left[\!\left[ \partial_{\mathbf{n}} w_h \right]\!\right]_{ij} \left[\left[\partial_{\mathbf{n}} \phi_h\right]\right]_{ij} \Big] \, \mathrm{d}\Gamma. \end{aligned} \tag{4.3}$$

For each interface facet $E \subset \Gamma_{ij}$, let $h_E$ denote the average diameter of the two adjacent elements. The weights in Equation (4.3) are taken facetwise as

$$\eta_w := p \, \frac{D_{\mathrm{eff}}}{h_E^3}, \qquad \eta_\theta := \frac{12^{p+1}}{p} \, \frac{D_{\mathrm{eff}}}{h_E}. \tag{4.4}$$

The powers of $h_E$ follow from dimensional consistency of the two interface energies. The effective rigidity $D_{\mathrm{eff}}$ has units of force times length. Since the displacement jump has units of length and $\mathrm{d}\Gamma$ has units of length, $D_{\mathrm{eff}} h_E^{-3} \left[\!\left[ w_h \right]\!\right]^2 \mathrm{d}\Gamma$ has units of energy. The normal-rotation jump is dimensionless, so $D_{\mathrm{eff}} h_E^{-1} \left[\!\left[ \partial_{\mathbf{n}} w_h \right]\!\right]^2 \mathrm{d}\Gamma$ has the same units. Degree- and mesh-dependent scaling of this type is standard in weak multi-patch coupling for Kirchhoff-type plate and shell discretizations [31, 32]. The factors $p$ and $12^{p+1}/p$ are the fixed degree-dependent prefactors used in the present implementation; they are not universal constants required by the general theory. The interface penalties weakly enforce displacement and normal-rotation transmission across physical panel seams. Moment

and shear transmission arise through the variational equilibrium of the coupled formulation rather than through separate strong constraints.

## 4.5 Discrete unilateral foundation term

Let $w_{h,i} = w_h|_{\Omega_i}$ and $\phi_{h,i} = \phi_h|_{\Omega_i}$. The discrete unilateral Winkler term is

$$g_h^{\mu}(w_h;\phi_h) := \sum_{i=1}^{M} \int_{\Omega_i} k_s \left[\chi_{\{w_{h,i}<0\}} + \varepsilon\, \chi_{\{w_{h,i}\geq 0\}}\right] w_{h,i}\phi_{h,i}\, \mathrm{d}\Omega, \qquad 0 < \varepsilon \ll 1. \tag{4.5}$$

The corresponding discrete load functional is

$$\ell_h^{\mu}(\phi_h) := \sum_{i=1}^{M} \int_{\Omega_i} q(x,y)\, \phi_{h,i}\, \mathrm{d}\Omega. \tag{4.6}$$

The foundation term retains the same active-set structure as the continuous model: its tangent stiffness depends on whether the current displacement is in compression or uplift [69, 70].

## 4.6 Mechanical full-order residual problem

Collecting the bending, interior-stabilization, interface, foundation, and load terms gives the finite-dimensional residual problem that defines the mechanical full-order model.

**Mechanical full-order residual problem**

Given $\boldsymbol{\mu} \in \mathcal{D}$, find $w_h(\boldsymbol{\mu}) \in V_h$ such that

$$\begin{aligned}\mathcal{F}_h^{\mu}(w_h;\phi_h) := a_h^{\mu}(w_h,\phi_h) + s_h^{\mu}(w_h,\phi_h) + c_{\Gamma,h}^{\mu}(w_h,\phi_h) \\ + g_h^{\mu}(w_h;\phi_h) - \ell_h^{\mu}(\phi_h) = 0 \qquad \forall\, \phi_h \in V_h.\end{aligned} \tag{4.7}$$

Within a fixed active-set branch, $w_h$ is equivalently a stationary point of the discrete total potential energy

$$\begin{aligned}\Pi_h^{\mu}(w_h) := \frac{1}{2}a_h^{\mu}(w_h,w_h) + \frac{1}{2}s_h^{\mu}(w_h,w_h) + \frac{1}{2}c_{\Gamma,h}^{\mu}(w_h,w_h) \\ + \frac{1}{2}g_h^{\mu}(w_h;w_h) - \ell_h^{\mu}(w_h).\end{aligned} \tag{4.8}$$

The solvability statement used throughout the parameterized study requires the orthotropic bending matrix to remain uniformly positive definite over the admissible parameter set. In the present notation, the corresponding conditions are

$$D_x > 0, \qquad D_y > 0, \qquad D_s > 0, \qquad D_x D_y - D_{xy}^2 > 0. \tag{4.9}$$

Together with sufficiently large interior-facet and physical-interface penalty coefficients, and with a uniformly positive foundation stiffness $k_s \geq k_0 > 0$, these conditions give a coercive stabilized bending contribution and a monotone regularized foundation response. Under these

assumptions, the finite-dimensional problem admits a unique solution for every admissible $\boldsymbol{\mu} \in \mathcal{D}$.

The active set is updated during the nonlinear solution process, so (4.8) should be read as the branchwise energy representation associated with the current compression–uplift partition. On each fixed branch, the foundation coefficient is constant and the energy is quadratic. The complete regularized problem is piecewise quadratic and continuously assembled through the displacement-dependent branch selection; it should not be interpreted as one globally smooth quadratic functional. The residual equation (4.7) is the high-fidelity model used in Chapter 5 to generate reference solutions and snapshot data for the mechanical reduced-order models.

## 4.7 Nonlinear solution strategy

The residual equation (4.7) is nonlinear because the unilateral foundation term switches between compression and uplift according to the sign of $w_h$. It is solved by Newton's method in active-set form [33, 71, 72]. For a fixed parameter value $\boldsymbol{\mu}$, the index $k$ counts Newton iterations. At iteration $k$, the current displacement determines the panel-wise active regions

$$\mathcal{A}_{i,-}^{(k)} := \{\, \mathbf{x} \in \Omega_i : w_{h,i}^{(k)}(\mathbf{x}) < 0 \,\}, \qquad \mathcal{A}_{i,+}^{(k)} := \{\, \mathbf{x} \in \Omega_i : w_{h,i}^{(k)}(\mathbf{x}) \geq 0 \,\}.$$

The compression–uplift partition is frozen while the tangent problem is assembled and solved, and it is recomputed from the updated displacement before the next iteration. This branchwise linearization is the active-set specialization of Newton's method for the regularized contact-type constitutive law.

Given $w_h^{(k)} \in V_h$, the Newton correction $\delta w_h^{(k)} \in V_h$ satisfies

$$\mathcal{J}_h^{\boldsymbol{\mu}}\big(w_h^{(k)}; \delta w_h^{(k)}, \phi_h\big) = -\mathcal{F}_h^{\boldsymbol{\mu}}\big(w_h^{(k)}; \phi_h\big) \qquad \forall\, \phi_h \in V_h, \tag{4.10}$$

where

$$\mathcal{J}_h^{\boldsymbol{\mu}}\big(w_h^{(k)}; \delta w_h, \phi_h\big) := \frac{\mathrm{d}}{\mathrm{d}\xi} \mathcal{F}_h^{\boldsymbol{\mu}}\big(w_h^{(k)} + \xi\, \delta w_h; \phi_h\big)\Big|_{\xi=0}. \tag{4.11}$$

The Jacobian decomposes as

$$\begin{aligned} \mathcal{J}_h^{\boldsymbol{\mu}}\big(w_h^{(k)}; \delta w_h, \phi_h\big) = a_h^{\boldsymbol{\mu}}(\delta w_h, \phi_h) + s_h^{\boldsymbol{\mu}}(\delta w_h, \phi_h) + c_{\Gamma,h}^{\boldsymbol{\mu}}(\delta w_h, \phi_h) \\ + \mathrm{d}g_h^{\boldsymbol{\mu}}\big(w_h^{(k)}; \delta w_h, \phi_h\big), \end{aligned} \tag{4.12}$$

with the active-branch derivative

$$\mathrm{d}g_h^{\boldsymbol{\mu}}(w_h; \delta w_h, \phi_h) := \sum_{i=1}^{M} \int_{\Omega_i} k_s \left[\chi_{\{w_{h,i}<0\}} + \varepsilon\, \chi_{\{w_{h,i}\geq 0\}}\right] \delta w_{h,i} \phi_{h,i}\, \mathrm{d}\Omega. \tag{4.13}$$

The displacement is updated by

$$w_h^{(k+1)} = w_h^{(k)} + \delta w_h^{(k)}. \tag{4.14}$$

Let $\{\phi_i\}_{i=1}^{\mathcal{N}_h}$ be a basis of $V_h$, where $\mathcal{N}_h := \dim(V_h)$. Writing

$$w_h^{(k)} = \sum_{i=1}^{\mathcal{N}_h} (\mathbf{w}_h^{(k)})_i \phi_i, \qquad \delta w_h^{(k)} = \sum_{i=1}^{\mathcal{N}_h} (\delta \mathbf{w}_h^{(k)})_i \phi_i,$$

the Newton step has the algebraic form

$$\mathbf{J}_h(\mathbf{w}_h^{(k)})\, \delta \mathbf{w}_h^{(k)} = -\mathbf{F}_h(\mathbf{w}_h^{(k)}), \tag{4.15}$$

where

$$[\mathbf{F}_h(\mathbf{w}_h^{(k)})]_i := \mathcal{F}_h^{\mu}(w_h^{(k)}; \phi_i), \qquad [\mathbf{J}_h(\mathbf{w}_h^{(k)})]_{ij} := \mathcal{J}_h^{\mu}(w_h^{(k)}; \phi_j, \phi_i). \tag{4.16}$$

Starting from $\mathbf{w}_h^{(0)}$, the implemented nonlinear stopping test is

$$\left\|\mathbf{F}_h(\mathbf{w}_h^{(k)})\right\|_{\ell^2} \le \max\left\{\varepsilon_{\mathrm{abs}}, \varepsilon_{\mathrm{rel}} \left\|\mathbf{F}_h(\mathbf{w}_h^{(0)})\right\|_{\ell^2}\right\}, \tag{4.17}$$

where $\varepsilon_{\mathrm{abs}} > 0$ and $\varepsilon_{\mathrm{rel}} > 0$ are the prescribed absolute and relative residual tolerances. Thus, convergence is accepted when the algebraic residual is either below the absolute floor or has been reduced sufficiently relative to its initial value. The active set used in the accepted residual is the one induced by the current iterate $\mathbf{w}_h^{(k)}$.

The matrix $\mathbf{J}_h$ inherits the locality of the discrete formulation. Panel-interior bending, consistency, symmetry, stabilization, and foundation contributions populate the sub-blocks associated with individual panels. Off-diagonal coupling between different panel components arises only from the displacement and normal-rotation terms on non-empty physical interfaces $\Gamma_{ij}$. The assembled system is therefore sparse, but it is not block diagonal once neighboring panels are coupled.

## 4.8 Verification against analytical and three-dimensional references

Before the full-order model is used to generate snapshot data in Chapter 5, its accuracy and convergence behavior are assessed through mechanical verification tests. These comparisons exercise the numerical formulation and its relation to analytical, three-dimensional, refined numerical, and monolithic references; they are not experimental physical validation of a particular flooring product or connector [35, 36, 73]. Unless stated otherwise, all computations reported in the verification studies use the uplift regularization parameter

$$\varepsilon = 10^{-3}.$$

The interior $C^0$-IPG stabilization and the panel-interface penalties are implemented as defined in Sections 4.3 and 4.4.

The first verification case examines the monolithic single-panel setting without internal interfaces. The benchmark is a simply supported rectangular plate

$$\Omega = [0, L] \times [0, W], \qquad L = 2.0\,\mathrm{m}, \quad W = 1.0\,\mathrm{m}.$$

The rigidity parameters are

$$D_x = 1.00 \times 10^4\,\mathrm{N\,m}, \qquad D_y = 1.00 \times 10^4\,\mathrm{N\,m},$$

$$D_{xy} = 2.00 \times 10^3\,\mathrm{N\,m}, \qquad D_s = 4.00 \times 10^3\,\mathrm{N\,m}.$$

The plate is subjected to the uniform transverse load

$$f = -1.00 \times 10^4\,\mathrm{N/m^2},$$

and rests on a Winkler foundation with stiffness

$$k_s = 1.00 \times 10^5\,\mathrm{N/m^3}.$$

Simply supported boundary conditions are imposed on all four edges. For the reported field comparison, the finite element approximation uses polynomial degree $p = 2$ on a triangular mesh with 10116 cells, resulting in a discrete space of dimension 20503.

**Table 4.1:** Monolithic simply supported verification setup used for the mechanical full-order model.

| Quantity | Value |
|---|---|
| Domain | $\Omega = [0, L] \times [0, W]$, $L = 2.0\,\mathrm{m}$, $W = 1.0\,\mathrm{m}$ |
| Boundary condition | Simply supported on all four edges |
| Rigidities | $D_x = D_y = 1.00 \times 10^4\,\mathrm{N\,m}$ |
| Coupling rigidity | $D_{xy} = 2.00 \times 10^3\,\mathrm{N\,m}$ |
| Torsional rigidity | $D_s = 4.00 \times 10^3\,\mathrm{N\,m}$ |
| Foundation stiffness | $k_s = 1.00 \times 10^5\,\mathrm{N/m^3}$ |
| Uniform load | $f = -1.00 \times 10^4\,\mathrm{N/m^2}$ |
| Thickness used for the 3D reference | $t = 5.00 \times 10^{-2}\,\mathrm{m}$ |
| Equivalent isotropic solid constants | $\nu = 0.20$, $E = 9.216 \times 10^8\,\mathrm{Pa}$ |
| Finite element space | $p = 2$, 10116 triangular cells, 20503 degrees of freedom |

The full-order displacement $w_{\mathrm{fom}}$ is compared with two references. The first is a finitely evaluated analytical Navier sine-series reference $w_{\mathrm{ext}}$ for a simply supported rectangular orthotropic Kirchhoff–Love plate on a linear Winkler foundation [12, 13]. In the controlled verification configuration, the unilateral foundation remains on the compression branch in the loaded region, so the reference reduces to the linear Winkler branch rather than a mixed active-set problem. The second reference is a three-dimensional small-strain solid elasticity model, from which the mid-plane vertical displacement $w_{\mathrm{sol}}$ is extracted. This comparison is used only as an independent consistency check for a slender plate.

For the Navier reference, the governing equation in the compression-active branch is

$$D_x\,\partial_{xxxx} w + 2(D_{xy} + 2D_s)\,\partial_{xxyy} w + D_y\,\partial_{yyyy} w + k_s w = f \qquad \text{in } \Omega. \tag{4.18}$$

The expansion is

$$w_{\mathrm{ext}}(x, y) = \sum_{m=1}^{\infty} \sum_{n=1}^{\infty} W_{mn} \sin\left(\frac{m\pi x}{L}\right) \sin\left(\frac{n\pi y}{W}\right), \tag{4.19}$$

with

$$\alpha_m = \frac{m\pi}{L}, \qquad \beta_n = \frac{n\pi}{W},$$

and

$$W_{mn} = \frac{Q_{mn}}{D_x\alpha_m^4 + 2(D_{xy} + 2D_s)\alpha_m^2\beta_n^2 + D_y\beta_n^4 + k_s}. \tag{4.20}$$

For a uniform load $q(x, y) = f$,

$$Q_{mn} = \frac{16f}{\pi^2 mn} \tag{4.21}$$

for odd $m$ and $n$, while $Q_{mn} = 0$ otherwise.

For the three-dimensional reference, the solid domain is

$$\Omega^{3D} = (0, L) \times (0, W) \times (0, t), \qquad t = 5.00 \times 10^{-2}\,\text{m}.$$

The thickness $t$ is not an additional parameter of the two-dimensional Kirchhoff–Love plate solve, which is specified directly through the bending rigidities $D_x$, $D_y$, $D_{xy}$, and $D_s$. It is needed only to construct the independent three-dimensional solid reference used for this consistency comparison.

Let $\mathbf{u} = (u_x, u_y, u_z)$ be the displacement field. The small-strain tensor and isotropic stress tensor are

$$\boldsymbol{\varepsilon}(\mathbf{u}) = \frac{1}{2}\left(\nabla\mathbf{u} + \nabla\mathbf{u}^T\right), \qquad \boldsymbol{\sigma}(\mathbf{u}) = \lambda\ \mathrm{tr}\big(\boldsymbol{\varepsilon}(\mathbf{u})\big)\mathbf{I} + 2\mu\,\boldsymbol{\varepsilon}(\mathbf{u}),$$

where $\lambda$ and $\mu$ are the Lamé constants [11, 13]. With $D = D_x = D_y$, the isotropic-compatible reference uses

$$\nu = \frac{D_{xy}}{D}, \qquad D_s = \frac{D(1-\nu)}{2}, \qquad E = \frac{12D(1-\nu^2)}{t^3}.$$

For the verification parameters above, this gives

$$\nu = 0.20, \qquad E = 9.216 \times 10^8\,\text{Pa}.$$

The scalar field compared with the plate solution is the mid-plane displacement

$$w_{\text{sol}}(x, y) = u_z(x, y, t/2).$$

The three-dimensional reference serves exclusively as an independent consistency check against the slender solid model.

Figure 4.1 summarizes the displacement comparison. The computed fields are visually indistinguishable, and the absolute errors are two to three orders of magnitude smaller than the peak displacement. The peak downward displacements are $-9.472$ mm for the FOM, $-9.486$ mm for the 3D solid reference, and $-9.477$ mm for the analytical solution. Thus, the FOM differs from the analytical and 3D solid references by approximately 0.05% and 0.15%, respectively, in minimum displacement. The corresponding raw spatial $L^2(\Omega)$ norms are $7.072967 \times 10^{-3}\,\text{m}^2$, $7.069999 \times 10^{-3}\,\text{m}^2$, and $7.077110 \times 10^{-3}\,\text{m}^2$, respectively, so the FOM-to-reference discrepancies remain below 0.06% at the global-norm level.

The bending and twisting moments recovered from the converged full-

| Solution | Min. displacement [mm] | $\|w\|_{L^2(\Omega)}$ $[10^{-3}\,\text{m}^2]$ |
|---|---|---|
| Full-order model $w_{\text{fom}}$ | $-9.472$ | 7.072967 |
| 3D solid reference $w_{\text{sol}}$ | $-9.486$ | 7.069999 |
| Analytical reference $w_{\text{ext}}$ | $-9.477$ | 7.077110 |

**Table 4.2:** Displacement comparison for the monolithic simply supported verification case.

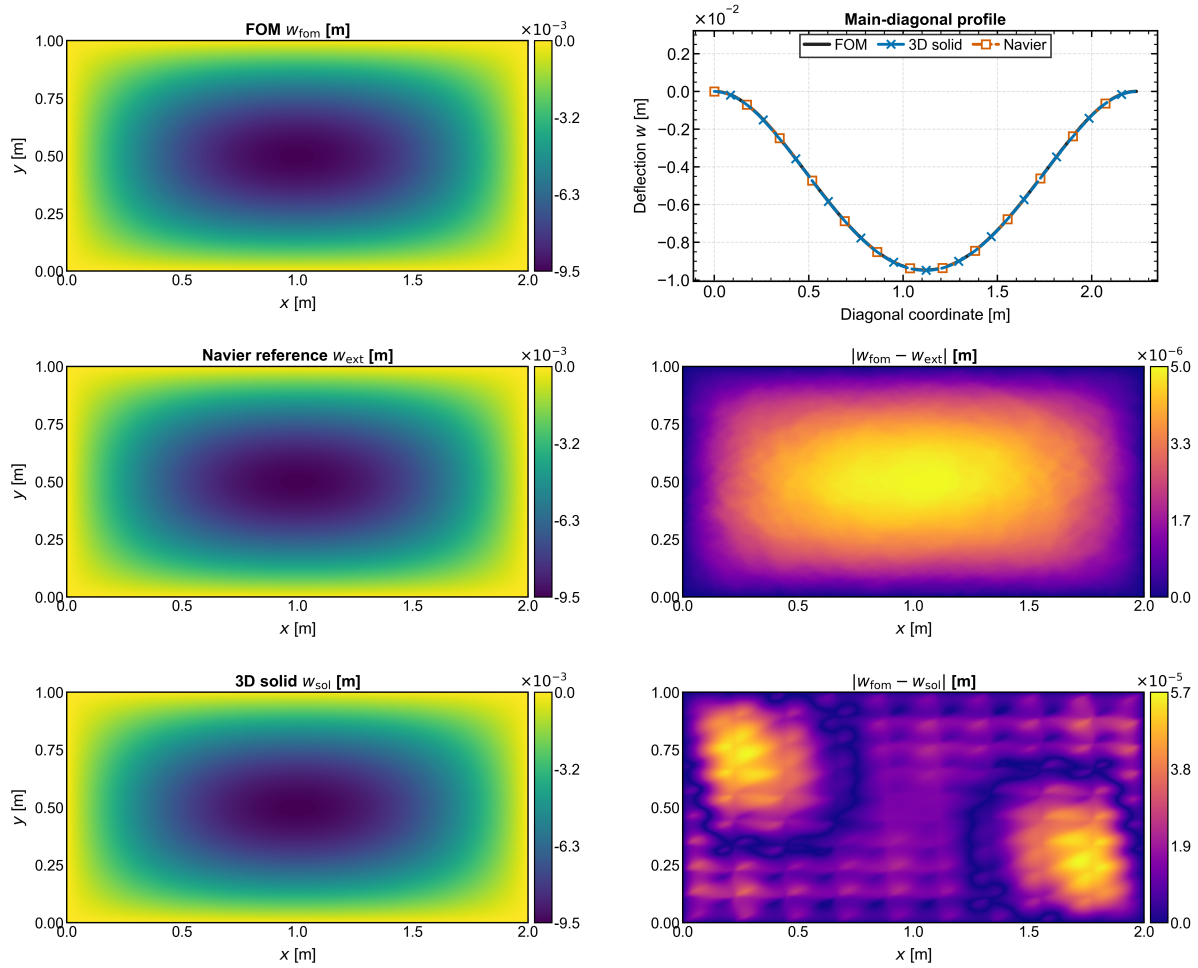



**Figure 4.1:** Verification summary for the monolithic simply supported plate under uniform load. Top row: full-order displacement $w_{\text{fom}}$ (left) and main-diagonal displacement profiles of the FOM, 3D solid reference, and analytical Navier-series reference (right). Middle row: analytical Navier-series displacement $w_{\text{ext}}$ (left) and absolute FOM–Navier error $|w_{\text{fom}} - w_{\text{ext}}|$ (right). Bottom row: 3D solid mid-plane displacement $w_{\text{sol}}$ (left) and absolute FOM–solid error $|w_{\text{fom}} - w_{\text{sol}}|$ (right).

order displacement are shown in Figure 4.2. The fields exhibit the expected qualitative behavior: the normal bending moments are consistent with the simply supported boundary response on the corresponding edges, and the twisting moment changes sign according to the symmetry of the deformation field.

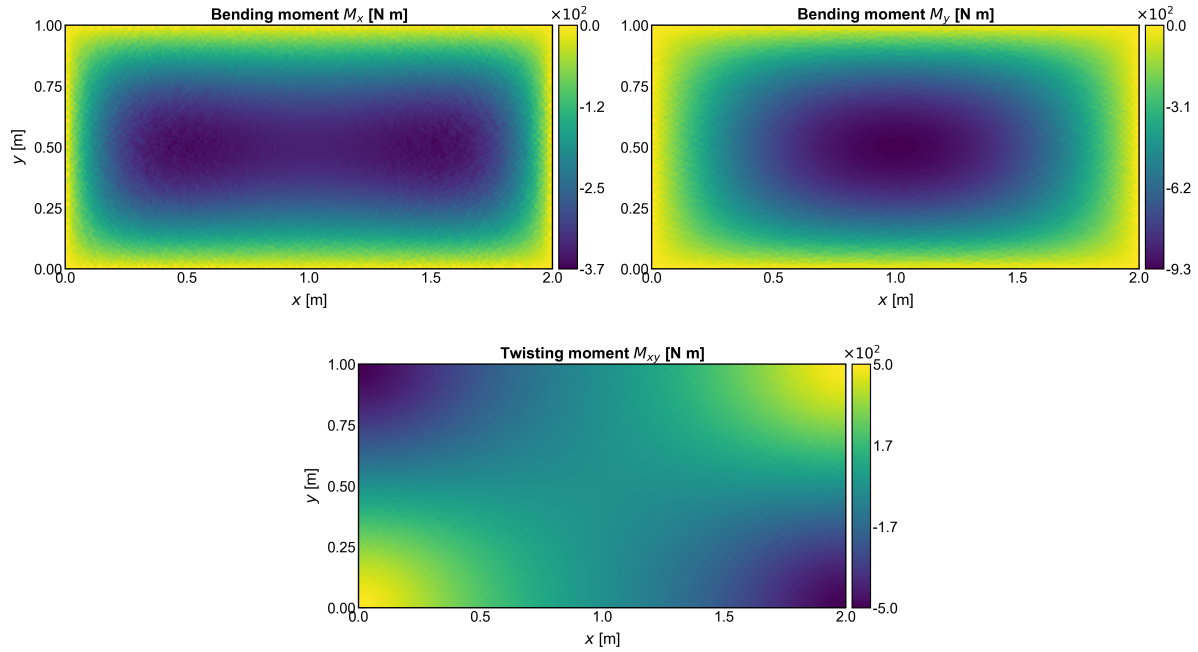



**Figure 4.2:** Moment-resultant fields recovered from the converged monolithic full-order displacement. Top row: bending moments $M_x$ (left) and $M_y$ (right). Bottom row: twisting moment $M_{xy}$, centered beneath the bending-moment fields.

## 4.9 Mesh-refinement convergence study

The spatial convergence of the monolithic discretization is assessed on five successively refined triangular meshes with nominal mesh lengths

$$\ell_{\mathcal{T}} \in \{0.5,\ 0.25,\ 0.125,\ 0.0625,\ 0.03125\}\,\mathrm{m}, \tag{4.22}$$

using polynomial degrees $p \in \{2, 3\}$. The resulting displacement spaces range from 37 to 11,548 degrees of freedom, thereby spanning approximately the $10^2$–$10^4$ range without referring to internal mesh-generator parameters. With the parameter ordering

$$\boldsymbol{\mu} = (D_x, D_y, D_{xy}, D_s, k_s, f),$$

the two material configurations are

$$\boldsymbol{\mu}_{\mathrm{iso}} = (10, 10, 3, 3.5, 1, -1) \cdot 10^3, \qquad \boldsymbol{\mu}_{\mathrm{ortho}} = (20, 10, 2.5, 3, 1, -1) \cdot 10^3, \tag{4.23}$$

with the rigidity, foundation, and load units established in Chapter 3.

Let $w_{\mathrm{ref}}$ denote the reference displacement and $w_h$ the finite element approximation. Before the errors are evaluated, both fields are represented in the same reference space. The dimensionless relative errors are

$$e_{L^2}(w_h) := \frac{\|w_{\mathrm{ref}} - w_h\|_{L^2(\Omega)}}{\|w_{\mathrm{ref}}\|_{L^2(\Omega)}}, \qquad e_{H^1}(w_h) := \frac{\|w_{\mathrm{ref}} - w_h\|_{H^1(\Omega)}}{\|w_{\mathrm{ref}}\|_{H^1(\Omega)}}, \tag{4.24}$$

where

$$\|v\|^2_{H^1(\Omega)} = \|v\|^2_{L^2(\Omega)} + \|\nabla v\|^2_{L^2(\Omega)}.$$

The reference norm is fixed within each configuration and therefore does not alter the asymptotic rate.

For smooth solutions, the plate-specific $C^0$ discontinuous-Galerkin estimates of Wells and Dung provide the relevant benchmark for triangular

Lagrange elements [27]; see also [25, 26]. In the present notation, the expected estimates are

$$\|w-w_h\|_{L^2(\Omega)} \lesssim \begin{cases} \ell_{\mathcal{T}}^2 \|w\|_{H^3(\Omega)}, & p=2, \\ \ell_{\mathcal{T}}^{p+1} \|w\|_{H^{p+1}(\Omega)}, & p \geq 3, \end{cases} \qquad \|w-w_h\|_{H^1(\Omega)} \lesssim \ell_{\mathcal{T}}^{p} \|w\|_{H^{p+1}(\Omega)}. \tag{4.25}$$

The corresponding expected orders are 2 in both $L^2$ and $H^1$ for $p = 2$, and 4 in $L^2$ together with 3 in $H^1$ for $p = 3$. The $H^1$ result is commonly stated for the seminorm. Here the displacement error has zero trace on the simply supported boundary, so the Poincaré inequality yields the same order for the full $H^1$ norm in Equation (4.24). The Winkler foundation contribution is of lower differential order and does not change this expected spatial hierarchy.

Two complementary reference constructions are used in Figure 4.3. In the first study, the linear Winkler plate approximation is compared with the analytical Navier reference for the simply supported plate under uniform loading. The series is evaluated using 30 odd indices in each coordinate direction and represented in a continuous degree-10 space on an $N_{\mathrm{ref}} = 64$ reference mesh. In the second study, the complete regularized unilateral FOM is compared with an independently computed fine-mesh FOM reference using $p_{\mathrm{ref}} = 4$ and $N_{\mathrm{ref}} = 64$, corresponding to 81,469 displacement degrees of freedom.

The two constructions answer different verification questions. The Navier comparison tests the discretization against an analytical solution on the linear foundation branch. The independent fine-mesh comparison retains the regularized unilateral formulation and tests the complete regularized unilateral mechanical implementation without using the same discrete solution as its reference. Agreement between the rates obtained from the two studies therefore provides stronger evidence than either comparison alone.

The rate triangles report the local slope between the two finest mesh levels,

$$r_{\mathrm{loc}}(e) := \frac{\log[e(\ell_{\mathcal{T},m-1})/e(\ell_{\mathcal{T},m})]}{\log[\ell_{\mathcal{T},m-1}/\ell_{\mathcal{T},m}]}, \qquad \ell_{\mathcal{T},m-1} = 0.0625\,\mathrm{m}, \quad \ell_{\mathcal{T},m} = 0.03125\,\mathrm{m}. \tag{4.26}$$

Against the analytical Navier reference, the measured $p = 2$ rates are 2.07–2.08 in $L^2$ and 2.05–2.07 in $H^1$. The corresponding $p = 3$ rates are 3.93–3.94 and 2.94–2.95, respectively, closely recovering the orders in Equation (4.25). The fine-mesh-reference study gives essentially the same $p = 2$ rates and $p = 3$ $H^1$ rates. For $p = 3$ in $L^2$, the local rates are 4.01 for the isotropic-compatible configuration and 3.72 for the orthotropic configuration. The slightly reduced orthotropic value is consistent with the finite numerical reference influencing the smallest-error level: the same orthotropic curve has rate 3.94 against the analytical Navier reference. In both configurations, the cubic approximation is uniformly more accurate than the quadratic one throughout the refinement sequence.

Taken together, the analytical and independent numerical studies verify the expected $L^2$ and $H^1$ convergence hierarchy of the monolithic $C^0$-IPG discretization for the tested Winkler-foundation configurations.

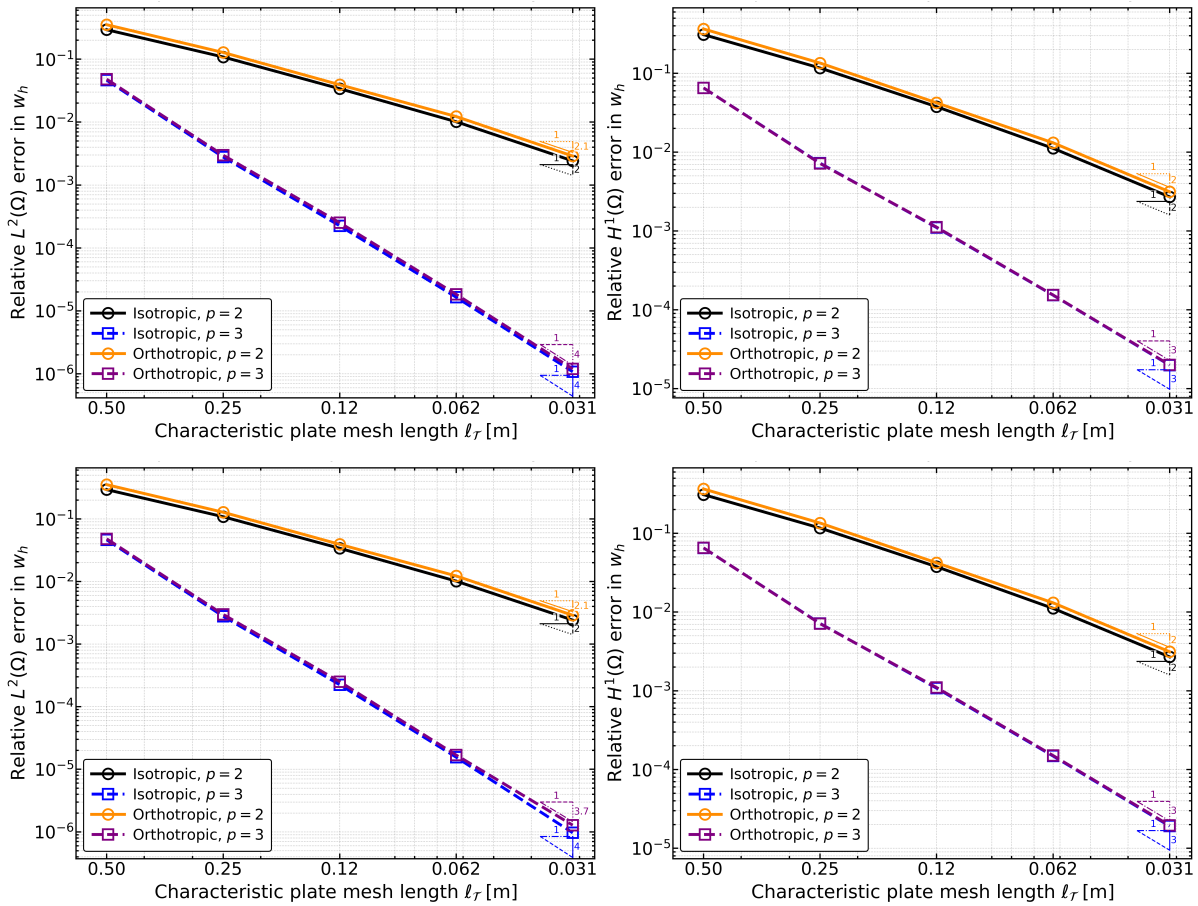



**Figure 4.3:** Relative displacement-error convergence for the monolithic simply supported plate under uniform loading with $k_s = 10^3\ \mathrm{N/m^3}$. Top: linear Winkler plate approximation against the analytical Navier reference. Bottom: regularized unilateral FOM approximation against an independent fine-mesh FOM reference ($p_{\mathrm{ref}} = 4$, $N_{\mathrm{ref}} = 64$). The left and right panels report the relative $L^2(\Omega)$ and $H^1(\Omega)$ errors, respectively, for the isotropic-compatible and orthotropic configurations with $p \in \{2, 3\}$. Triangles indicate local convergence rates computed from the two finest mesh levels.

## 4.10 Penalty-coupled panel-partition assessment

The panel-partition verification assesses the penalty-coupled formulation for an orthotropic plate subjected to a localized patch load. The goal is to compare the panel-wise formulation with a monolithic reference solution and to examine how the approximation behaves as the number of panel interfaces increases.

The orthotropic parameter vector is

$$\boldsymbol{\mu} = (20, 10, 2.5, 3, 1, -1) \cdot 10^3.$$

A localized patch load is centered at $(1.0, 0.5)$ m. Four partition layouts are considered, consisting of 2, 3, 4, and 6 panels, as shown in Figure 4.4. The polynomial degree is varied over $p \in \{2, 3\}$, and the interface weights are those defined in Section 4.4. The comparison therefore assesses the adopted discretization as a whole, including the combined effect of local polynomial degree and the corresponding interface weighting.

The performance of the panel-partition approximation is assessed through spatial error fields and relative $L^2$ errors. Figure 4.5 shows the absolute difference between each penalty-coupled panel-partition solution and the monolithic reference for $p = 2$. In every configuration, the error

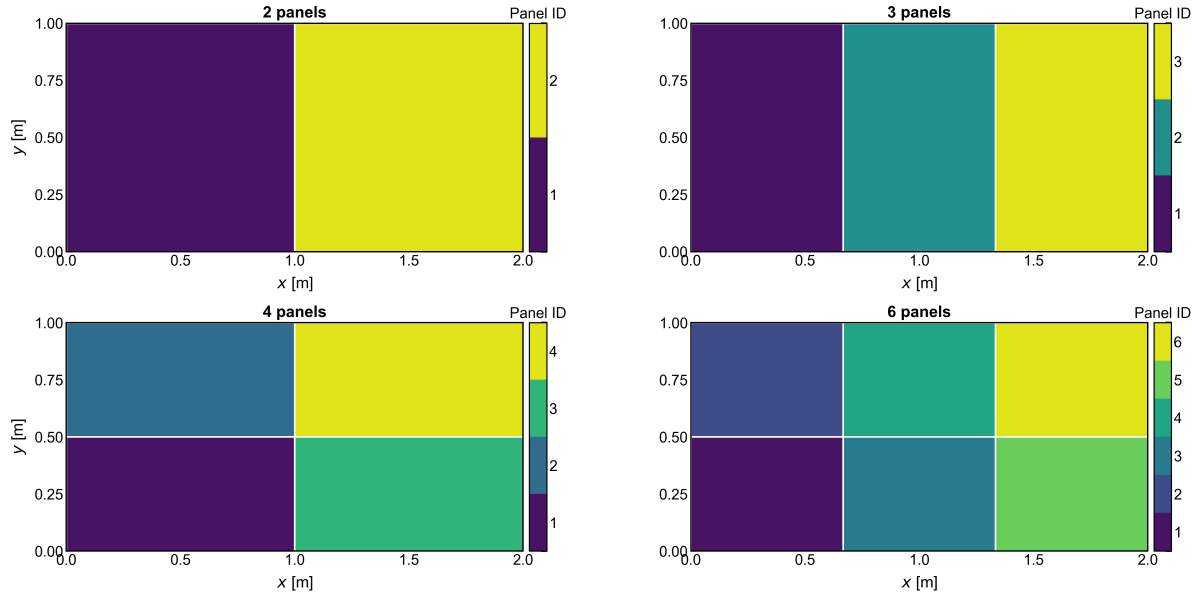



**Figure 4.4:** Panel-partition layouts considered for the orthotropic plate benchmark. Top row: two-panel configuration (left) and three-panel configuration (right). Bottom row: four-panel configuration (left) and six-panel configuration (right). The colors and corresponding panel-ID scales identify the individual subdomains and their internal interfaces. These configurations are used to assess the influence of weakly enforced panel interfaces on the full-order approximation.

is concentrated near the introduced panel interfaces, as expected for a weakly coupled formulation. The error remains small relative to the overall displacement, but it increases with the number of interfaces. In particular, the maximum error in the 6-panel case is about twice that of the 2-panel case, reflecting the cumulative effect of multiple weak interface couplings.

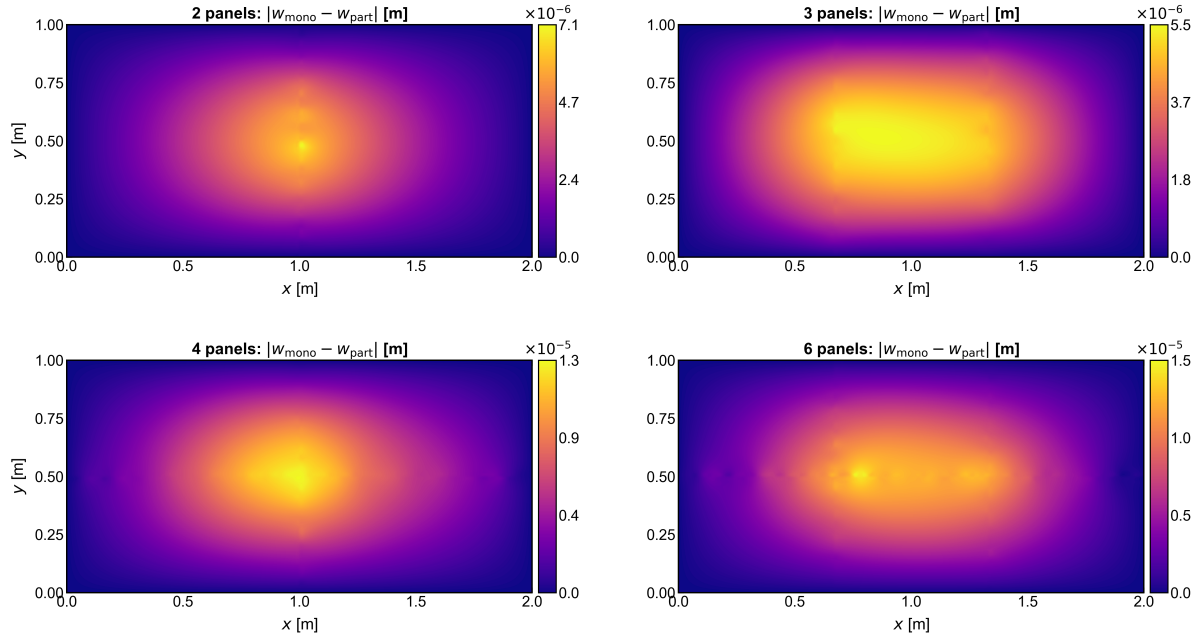



**Figure 4.5:** Absolute displacement-error fields $|w_{\text{mono}} - w_{\text{part}}|$ for the penalty-coupled panel configurations, computed with respect to the monolithic reference solution for $p = 2$. Top row: two-panel configuration (left) and three-panel configuration (right). Bottom row: four-panel configuration (left) and six-panel configuration (right). The discrepancies are concentrated primarily around the loaded central region and exhibit features associated with the introduced panel interfaces. Each panel uses its displayed color scale, with errors reported in meters.

The quantitative comparison in Figure 4.6 confirms the same trend. For each panel layout, increasing the polynomial degree from $p = 2$ to $p = 3$ reduces the relative $L^2$ error. This improvement reflects both a more accurate local approximation inside each panel and stronger effective enforcement of relative normal-rotation continuity across interfaces through the degree-dependent scaling. At the same time, the error increases with the number of panels, highlighting the cumulative contribution of the weakly enforced interfaces to the global approximation error. Within the tested resolution and penalty prescription, these experiments show that the panel-partition formulation remains close to the monolithic reference while retaining a modular panel-wise discretization.

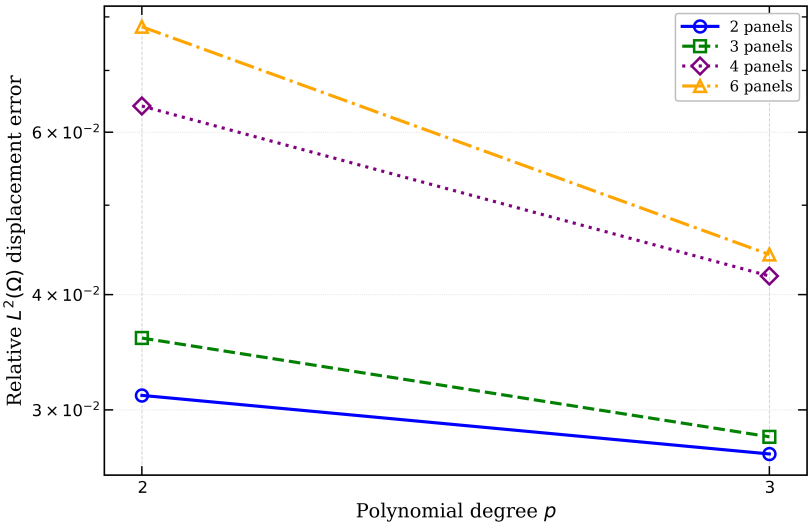



**Figure 4.6:** Relative $L^2(\Omega)$ displacement error with respect to the monolithic reference for the two-, three-, four-, and six-panel orthotropic configurations, reported for polynomial degrees $p \in \{2, 3\}$.

## 4.11 Summary of the full-order model

The mechanical high-fidelity model uses standard $H^1$-conforming Lagrange spaces augmented by a symmetric $C^0$ interior-penalty treatment on numerical facets strictly inside each physical panel. Displacement and normal-rotation transmission across physical seams are assembled through separate penalties whose mesh powers follow dimensional consistency and whose degree-dependent prefactors are implementation choices. The regularized unilateral Winkler law remains the only nonlinear contribution and is solved by an active-set Newton method with the algebraic residual stopping test in Equation (4.17).

The verification study tests complementary aspects of this construction. In the monolithic simply supported benchmark, the displacement closely matches the finitely evaluated Navier-series and independent three-dimensional solid references, while the recovered moments exhibit the expected symmetry and boundary behavior. Across five refinement levels, the $p = 2$ and $p = 3$ solutions recover the plate-specific $L^2$–$H^1$ convergence hierarchy against both analytical Navier and fine-mesh FOM references. The panel-partition study shows that finite interface penalties concentrate the dominant discrepancies near physical seams, that errors accumulate with additional interfaces, and that higher polynomial degree with order-aware scaling reduces these errors. Within the verification scope of Section 4.8, Equation (4.7) defines the mechanical high-fidelity model used in Chapter 5 for snapshot generation, intrusive projection, and reduced-order assessment.

# Reduced-Order Modeling of the Mechanical Plate Problem

# 5


5.1 Many-query setting and reduced-order objective . . 42
5.2 POD basis construction and intrusive POD–Galerkin reduction . . . . . 42
5.3 Non-intrusive POD-coordinate surrogates . . . 44
5.4 Assessment protocol and offline–online accounting . 48
5.5 ROM assessment for high-dimensional mechanical parametrization . . . . . 49
5.6 ROM assessment for penalty-coupled panel systems . . . . . 52
5.7 ROM assessment for multi-panel configurations . . . . 56
5.8 Summary and limitations of the mechanical ROMs . . . 59


## 5.1 Many-query setting and reduced-order objective

Chapters 3 and 4 define the parameterized mechanical plate model and its high-fidelity finite element approximation. The full-order model provides the reference solution throughout the admissible parameter domain, but repeated assembly and solution of its nonlinear algebraic system remain expensive in many-query settings. A reduced-order model replaces the full-order trial space by a problem-adapted space of substantially smaller dimension and represents each new displacement field through a limited number of reduced coordinates [74–76]. The computational work is therefore divided into an *offline stage*, in which representative high-fidelity solutions are generated and compressed, and an *online stage*, in which the reduced coordinates are determined for a new parameter value.

The reduced-order objective is to approximate, for many $\boldsymbol{\mu} \in \mathcal{D}$, the high-fidelity displacement $w_h(\boldsymbol{\mu}) \in V_h$ without modifying the parent mechanical model. Proper Orthogonal Decomposition (POD) is used to construct the reduced spaces. Two online strategies are considered. The intrusive POD–Galerkin model, abbreviated POD–G, projects the full-order residual and tangent operator onto a reduced trial space. The non-intrusive POD-coordinate surrogates instead learn the map from the physical parameters to the retained POD coordinates by interpolation, Gaussian-process regression, or a neural network. Both strategies start from high-fidelity snapshots, but they differ in how the reduced coordinates are evaluated and in the amount of full-order information retained online [39–41, 43, 46–48, 77].

## 5.2 POD basis construction and intrusive POD–Galerkin reduction

Let $\mathcal{N}_h := \dim(V_h)$, and let

$$\mathbf{w}_h(\boldsymbol{\mu}) \in \mathbb{R}^{\mathcal{N}_h}$$

denote the coefficient vector of the high-fidelity displacement in the finite element basis introduced in Chapter 4. The discrete solution manifold is

$$\mathcal{M}_h := \{\, \mathbf{w}_h(\boldsymbol{\mu}) : \boldsymbol{\mu} \in \mathcal{D} \,\} \subset \mathbb{R}^{\mathcal{N}_h}. \tag{5.1}$$

A set of $N_s$ offline parameters

$$\{\boldsymbol{\mu}^{(i)}\}_{i=1}^{N_s} \subset \mathcal{D}$$

is selected by Latin hypercube sampling [78, 79]. For each parameter, the nonlinear full-order problem is solved by the active-set Newton strategy of Section 4.7, and the corresponding coefficient vectors are collected in

the snapshot matrix

$$\mathbf{S} = \left[\mathbf{w}_h(\boldsymbol{\mu}^{(1)}) \,\middle|\, \cdots \,\middle|\, \mathbf{w}_h(\boldsymbol{\mu}^{(N_s)})\right] \in \mathbb{R}^{\mathcal{N}_h \times N_s}. \tag{5.2}$$

The POD metric is the discrete counterpart of the displacement-plus-gradient inner product used in the implementation. For panel-wise fields $u_h = \{u_{h,m}\}_{m=1}^{M}$ and $v_h = \{v_{h,m}\}_{m=1}^{M}$, it is

$$(u_h, v_h)_{\mathrm{POD},h} := \sum_{m=1}^{M} \int_{\Omega_m} \left(u_{h,m} v_{h,m} + \nabla u_{h,m} \cdot \nabla v_{h,m}\right) \mathrm{d}\Omega. \tag{5.3}$$

For a monolithic plate, the sum contains one term. If $\{\phi_a\}_{a=1}^{\mathcal{N}_h}$ denotes the full-order basis, the corresponding Gram matrix is

$$[\mathbf{W}_{\mathrm{POD}}]_{ab} := (\phi_b, \phi_a)_{\mathrm{POD},h}, \qquad \mathbf{W}_{\mathrm{POD}} \in \mathbb{R}^{\mathcal{N}_h \times \mathcal{N}_h}. \tag{5.4}$$

The displacement term in Equation (5.3) makes the metric positive definite on the discrete space.

The weighted method of snapshots forms the correlation matrix [53]

$$\mathbf{C} := \mathbf{S}^{\top} \mathbf{W}_{\mathrm{POD}} \mathbf{S} \in \mathbb{R}^{N_s \times N_s}. \tag{5.5}$$

Let $r := \operatorname{rank}(\mathbf{S})$, and let $\{(\lambda_j, \mathbf{q}_j)\}_{j=1}^{r}$ be the nonzero eigenpairs of $\mathbf{C}$, with Euclidean-orthonormal eigenvectors and

$$\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_r > 0.$$

For a retained dimension $N$ satisfying $1 \leq N \leq r$, the POD modes and basis matrix are

$$\boldsymbol{\psi}_j := \lambda_j^{-1/2} \mathbf{S} \mathbf{q}_j, \quad j = 1, \ldots, N, \qquad \mathbf{V} := [\boldsymbol{\psi}_1 \,|\, \cdots \,|\, \boldsymbol{\psi}_N] \in \mathbb{R}^{\mathcal{N}_h \times N}. \tag{5.6}$$

The retained modes satisfy

$$\mathbf{V}^{\top} \mathbf{W}_{\mathrm{POD}} \mathbf{V} = \mathbf{I}_N. \tag{5.7}$$

Because the eigenvalues are ordered decreasingly, the first modes identify the dominant directions of the uncentered snapshot ensemble in the POD metric. Among all $N$-dimensional subspaces of the snapshot span, their span minimizes the total squared metric projection error over the snapshots, and the minimum is the discarded eigenvalue sum $\sum_{j=N+1}^{r} \lambda_j$. The cumulative retained-energy ratio is therefore

$$E_N := \frac{\sum_{j=1}^{N} \lambda_j}{\sum_{j=1}^{r} \lambda_j}, \tag{5.8}$$

and the nominal retained dimension is selected from

$$N := \min\{\, n \in \{1, \ldots, r\} : E_n \geq \eta_{\mathrm{POD}} \,\}, \qquad 0 < \eta_{\mathrm{POD}} < 1. \tag{5.9}$$

The value of $\eta_{\mathrm{POD}}$, together with the retained-dimension selection procedure, is specified for each numerical test.

For the intrusive model, the reduced approximation is sought directly in the high-fidelity trial space as

$$\mathbf{w}_N(\boldsymbol{\mu}) = \mathbf{V} \mathbf{z}(\boldsymbol{\mu}), \qquad \mathbf{z}(\boldsymbol{\mu}) \in \mathbb{R}^N. \tag{5.10}$$

No projection of a known high-fidelity state is required online; the coordinates are determined by the reduced equilibrium equations. Let $\mathbf{F}_h(\mathbf{w}_h;\boldsymbol{\mu})$ denote the algebraic residual defined from Section 4.6, and define

$$\mathbf{J}_h(\mathbf{w}_h;\boldsymbol{\mu}) := \frac{\partial \mathbf{F}_h(\mathbf{w}_h;\boldsymbol{\mu})}{\partial \mathbf{w}_h}.$$

The reduced residual and tangent matrix are

$$\mathbf{F}_N(\mathbf{z};\boldsymbol{\mu}) := \mathbf{V}^\top \mathbf{F}_h(\mathbf{V}\mathbf{z};\boldsymbol{\mu}), \qquad \mathbf{J}_N(\mathbf{z};\boldsymbol{\mu}) := \mathbf{V}^\top \mathbf{J}_h(\mathbf{V}\mathbf{z};\boldsymbol{\mu})\mathbf{V}. \tag{5.11}$$

The intrusive POD–Galerkin problem is

**Intrusive mechanical POD–Galerkin model**

$$\mathbf{F}_N(\mathbf{z};\boldsymbol{\mu}) = \mathbf{0}. \tag{5.12}$$

At iteration $k$ of the reduced Newton method for a fixed $\boldsymbol{\mu}$, the correction satisfies

$$\mathbf{J}_N(\mathbf{z}^{(k)};\boldsymbol{\mu})\,\delta\mathbf{z}^{(k)} = -\mathbf{F}_N(\mathbf{z}^{(k)};\boldsymbol{\mu}), \qquad \mathbf{z}^{(k+1)} = \mathbf{z}^{(k)} + \delta\mathbf{z}^{(k)}. \tag{5.13}$$

The compression–uplift active set is evaluated from the lifted iterate $\mathbf{V}\mathbf{z}^{(k)}$ [80]. Since the snapshots satisfy the same homogeneous essential boundary conditions, every lifted reduced iterate remains admissible. In the present implementation, the nonlinear residual and Jacobian are still assembled in the full-order space before projection, whereas only the linear solve is reduced to dimension $N$. The online cost is therefore not independent of $\mathcal{N}_h$, which limits the speedup of this non-hyper-reduced intrusive realization.

## 5.3 Non-intrusive POD-coordinate surrogates

Non-intrusive POD-coordinate surrogates determine the reduced coordinates without assembling or projecting the governing residual online [46, 47]. For a panel-partitioned full-order state $w_h = \{w_{h,m}\}_{m=1}^{M}$, the coefficient vector belongs to a broken product space and depends on the panel-component ordering. The learned parameter-to-coordinate map is instead constructed from a single physical displacement field represented in a fixed global continuous space $V_{\mathrm{CG}} \subset H^1(\Omega)$. The transferred field $\widetilde{w}_h \in V_{\mathrm{CG}}$ is defined by the $L^2$-projection

$$\int_\Omega \widetilde{w}_h\, v_{\mathrm{CG}}\,\mathrm{d}\Omega = \sum_{m=1}^{M} \int_{\Omega_m} w_{h,m}\, v_{\mathrm{CG}}\,\mathrm{d}\Omega \qquad \forall v_{\mathrm{CG}} \in V_{\mathrm{CG}}. \tag{5.14}$$

The right-hand side represents the assembled panel-wise physical field almost everywhere, while the projection supplies a continuous coefficient vector with a parameter-independent ordering. The parent state remains the broken FOM field $w_h = \{w_{h,m}\}_{m=1}^{M}$; the transferred field $\widetilde{w}_h$ is the output representation used for non-intrusive learning, reconstruction, and common field comparison. For a monolithic plate, the same operation is a compatible transfer to $V_{\mathrm{CG}}$. The transfer does not modify the FOM or the intrusive POD–Galerkin equations. Because $V_{\mathrm{CG}}$ is continuous, displacement jumps and normal-rotation jumps at physical seams are not retained as separate ROM outputs. An interface-aware surrogate for

those quantities would require either additional output channels or a reduced representation of the broken panel state.

Let $\{\chi_a\}_{a=1}^{\mathcal{N}_{\mathrm{CG}}}$ be a basis of $V_{\mathrm{CG}}$, where $\mathcal{N}_{\mathrm{CG}} := \dim(V_{\mathrm{CG}})$, and let $\widetilde{\mathbf{w}}_{\mathrm{CG}}(\boldsymbol{\mu}^{(i)})$ denote the coefficient vector of the transferred snapshot. The transferred snapshot matrix is

$$\mathbf{S}_{\mathrm{CG}} := \left[\widetilde{\mathbf{w}}_{\mathrm{CG}}(\boldsymbol{\mu}^{(1)}) \,\middle|\, \cdots \,\middle|\, \widetilde{\mathbf{w}}_{\mathrm{CG}}(\boldsymbol{\mu}^{(N_s)})\right] \in \mathbb{R}^{\mathcal{N}_{\mathrm{CG}} \times N_s}. \tag{5.15}$$

The corresponding global POD metric and Gram matrix are

$$(u, v)_{\mathrm{POD,CG}} := \int_\Omega (uv + \nabla u \cdot \nabla v)\, \mathrm{d}\Omega, \qquad [\mathbf{W}_{\mathrm{POD}}^{\mathrm{CG}}]_{ab} := (\chi_b, \chi_a)_{\mathrm{POD,CG}}. \tag{5.16}$$

Applying Equations (5.5)–(5.9) with $\mathbf{S}_{\mathrm{CG}}$ and $\mathbf{W}_{\mathrm{POD}}^{\mathrm{CG}}$ gives the global basis

$$\mathbf{V}_{\mathrm{CG}} = [\boldsymbol{\psi}_1^{\mathrm{CG}} \,|\, \cdots \,|\, \boldsymbol{\psi}_N^{\mathrm{CG}}].$$

For each offline snapshot, the target coordinate vector is its metric projection onto this basis:

$$\mathbf{G}_{\mathrm{CG}}\mathbf{z}^{(i)} = \mathbf{V}_{\mathrm{CG}}^{\top}\mathbf{W}_{\mathrm{POD}}^{\mathrm{CG}}\widetilde{\mathbf{w}}_{\mathrm{CG}}(\boldsymbol{\mu}^{(i)}), \qquad \mathbf{G}_{\mathrm{CG}} := \mathbf{V}_{\mathrm{CG}}^{\top}\mathbf{W}_{\mathrm{POD}}^{\mathrm{CG}}\mathbf{V}_{\mathrm{CG}} = \mathbf{I}_N. \tag{5.17}$$

The resulting regression data are $\{(\boldsymbol{\mu}^{(i)}, \mathbf{z}^{(i)})\}_{i=1}^{N_s}$. Parameters and coordinates are standardized with component-wise affine maps fitted exclusively on the regression-training index set $\mathcal{I}_{\mathrm{tr}}$:

$$\widehat{\boldsymbol{\mu}} = \mathbf{D}_\mu(\boldsymbol{\mu} - \overline{\boldsymbol{\mu}}), \qquad \widehat{\mathbf{z}} = \mathbf{D}_z(\mathbf{z} - \overline{\mathbf{z}}). \tag{5.18}$$

Here $\overline{\boldsymbol{\mu}}$ and $\overline{\mathbf{z}}$ are component-wise means over $\mathcal{I}_{\mathrm{tr}}$, while

$$\mathbf{D}_\mu = \mathrm{diag}(s_{\mu,1}^{-1}, \ldots, s_{\mu,P}^{-1}), \qquad \mathbf{D}_z = \mathrm{diag}(s_{z,1}^{-1}, \ldots, s_{z,N}^{-1})$$

contain the reciprocal component-wise standard deviations computed from the same subset; a zero-variance component is assigned unit scale. These affine transformations apply only to the regression variables and do not alter the uncentered POD snapshots. Once a surrogate predicts $\widehat{\mathbf{z}}(\widehat{\boldsymbol{\mu}})$, the reconstructed displacement is

$$\widetilde{\mathbf{w}}_N(\boldsymbol{\mu}) = \mathbf{V}_{\mathrm{CG}}\left(\overline{\mathbf{z}} + \mathbf{D}_z^{-1}\widehat{\mathbf{z}}(\widehat{\boldsymbol{\mu}})\right). \tag{5.19}$$

Relative to the transferred high-fidelity field, the non-intrusive discrepancy contains the finite-basis approximation error and the error of the learned coordinate map. The direct projection in Equation (5.17), denoted POD–Proj in the numerical comparisons, provides the retained-basis projection baseline. Since it requires an already available high-fidelity solution, POD–Proj is a diagnostic and not a deployable online ROM.

## POD interpolation: PODI–RBF and PODI–Linear

The PODI methods approximate the map from standardized parameters to standardized coordinates [46]. Their regression-training data are

$$\left\{\left(\widehat{\boldsymbol{\mu}}^{(i)}, \widehat{\mathbf{z}}^{(i)}\right)\right\}_{i \in \mathcal{I}_{\mathrm{tr}}}.$$

For PODI–RBF, the implementation uses a vector-valued thin-plate-spline interpolant with zero smoothing in the standardized parameter

and coordinate variables [81, 82]:

$$\widehat{\mathbf{z}}(\widehat{\boldsymbol{\mu}}) = \sum_{i \in \mathcal{I}_{\mathrm{tr}}} \boldsymbol{\alpha}_i \, \varphi\left(\left\|\widehat{\boldsymbol{\mu}} - \widehat{\boldsymbol{\mu}}^{(i)}\right\|_2\right) + \mathbf{b}_0 + \mathbf{B}\widehat{\boldsymbol{\mu}}, \qquad \varphi(\rho) = \begin{cases} \rho^2 \log \rho, & \rho > 0, \\ 0, & \rho = 0. \end{cases} \tag{5.20}$$

Here $\boldsymbol{\alpha}_i, \mathbf{b}_0 \in \mathbb{R}^N$ and $\mathbf{B} \in \mathbb{R}^{N \times P}$. The polynomial degree is left at the kernel-dependent default rather than selected empirically; for the thin-plate-spline kernel this gives the degree-one affine term $\mathbf{b}_0 + \mathbf{B}\widehat{\boldsymbol{\mu}}$. The interpolation coefficients satisfy the standard zeroth- and first-moment side constraints

$$\sum_{i \in \mathcal{I}_{\mathrm{tr}}} \boldsymbol{\alpha}_i = \mathbf{0}, \qquad \sum_{i \in \mathcal{I}_{\mathrm{tr}}} \widehat{\mu}_\ell^{(i)} \boldsymbol{\alpha}_i = \mathbf{0}, \quad \ell = 1, \ldots, P. \tag{5.21}$$

For PODI–Linear, the standardized training parameters form the vertices of a Delaunay tessellation in the $P$-dimensional physical-parameter space [83, 84]. This simplicial construction is unrelated to the physical finite element mesh. Standardization is required because otherwise differences in units or numerical ranges would distort the Euclidean geometry used to construct and search the simplices. Let a standardized query $\widehat{\boldsymbol{\mu}}$ lie in the containing simplex

$$\mathcal{S} = \operatorname{conv}\left\{\widehat{\boldsymbol{\mu}}^{(i_0)}, \ldots, \widehat{\boldsymbol{\mu}}^{(i_P)}\right\}, \qquad i_0, \ldots, i_P \in \mathcal{I}_{\mathrm{tr}}.$$

Its barycentric weights $\beta_0, \ldots, \beta_P$ and the interpolated coordinate vector satisfy

$$\sum_{s=0}^{P} \beta_s \widehat{\boldsymbol{\mu}}^{(i_s)} = \widehat{\boldsymbol{\mu}}, \qquad \sum_{s=0}^{P} \beta_s = 1, \qquad \beta_s \geq 0, \qquad \widehat{\mathbf{z}}(\widehat{\boldsymbol{\mu}}) = \sum_{s=0}^{P} \beta_s \widehat{\mathbf{z}}^{(i_s)}. \tag{5.22}$$

A containing simplex exists only inside the convex hull of the standardized training parameters. If a query lies outside that convex hull, or if the simplicial interpolator returns a non-finite coordinate vector, the implementation selects

$$i_\star \in \arg\min_{i \in \mathcal{I}_{\mathrm{tr}}} \left\|\widehat{\boldsymbol{\mu}} - \widehat{\boldsymbol{\mu}}^{(i)}\right\|_2$$

and returns $\widehat{\mathbf{z}}^{(i_\star)}$ as the deterministic nearest-training-neighbor fallback.

## POD neural-network regression

The POD–NN model follows the non-intrusive parameter-to-coordinate strategy of Hesthaven and Ubbiali [47]. It is a fully connected regressor

$$\Phi_\theta : \mathbb{R}^P \longrightarrow \mathbb{R}^N$$

with hidden widths 128, 96, and 64, ELU hidden activations, and a linear $N$-dimensional output. The executed mechanical network uses neither dropout nor layer normalization. The regression records are partitioned into training, validation, and internal-holdout subsets using an 80/10/10 split. Parameter and coordinate standardizers are fitted on the training subset only, and the target vectors are the precomputed standardized

POD projection coordinates. The coefficient-only training objective is

$$\mathscr{L}_{\mathrm{coef}}(\theta) = \frac{1}{|\mathcal{I}_{\mathrm{tr}}|N} \sum_{i \in \mathcal{I}_{\mathrm{tr}}} \left\| \widehat{\mathbf{z}}^{(i)} - \Phi_\theta\big(\widehat{\boldsymbol{\mu}}^{(i)}\big) \right\|_2^2 . \tag{5.23}$$

Training uses AdamW with batch size 256, initial learning rate $5 \times 10^{-4}$, weight decay $2 \times 10^{-6}$, and at most 2500 epochs. The validation coefficient loss controls learning-rate reduction, early stopping, and checkpoint selection; after training, the state attaining the minimum validation loss is restored. The internal-holdout subset is used only for diagnostic evaluation and does not enter optimization or model selection.

The POD projection coordinates are computed once during offline data preparation. For a new parameter, the predicted standardized coordinates are inverse-transformed and the displacement is reconstructed once through the fixed POD basis according to Equation (5.19).

## POD Gaussian-process regression

Gaussian-process regression supplies a coordinate prediction and a model-based posterior uncertainty estimate [48, 77]. One independent zero-mean Gaussian process is fitted to each standardized POD coordinate. For $j = 1, \ldots, N$,

$$\widehat{z}_j(\widehat{\boldsymbol{\mu}}) = g_j(\widehat{\boldsymbol{\mu}}) + \epsilon_j, \qquad g_j \sim \mathcal{GP}(0, k_j^{(0)}), \qquad \epsilon_j \sim \mathcal{N}(0, \sigma_{n,j}^2).$$

The smooth covariance component is the constant-amplitude anisotropic squared-exponential kernel

$$k_j^{(0)}(\widehat{\boldsymbol{\mu}}, \widehat{\boldsymbol{\mu}}') = \sigma_{f,j}^2 \exp\left[ -\frac{1}{2} \sum_{\ell=1}^{P} \frac{(\widehat{\mu}_\ell - \widehat{\mu}'_\ell)^2}{\ell_{j,\ell}^2} \right], \tag{5.24}$$

where $\sigma_{f,j}^2 > 0$ is the signal variance and $\ell_{j,\ell} > 0$ are automatic-relevance-determination length scales. The fitted white-noise variance $\sigma_{n,j}^2 \geq 0$ is included when that kernel component is enabled and is set to zero otherwise. A separate small positive numerical nugget $\alpha_j$ is added to the covariance diagonal for stabilization.

The POD–GPR data use the same 80/10/10 training/validation/internal-holdout roles as the mechanical POD–NN. The parameter and coordinate scalers and all Gaussian-process hyperparameters are fitted from the training subset only. Validation and internal-holdout data are diagnostic and do not enter the fit. No additional target normalization is applied inside the Gaussian-process estimator because the POD coordinates have already been standardized. The enabled kernel hyperparameters are estimated by log-marginal-likelihood optimization with multiple optimizer starts.

Let $\widehat{\mathbf{z}}_{j,\mathrm{tr}} = (\widehat{z}_j^{(a)})_{a \in \mathcal{I}_{\mathrm{tr}}}$. For $a, b \in \mathcal{I}_{\mathrm{tr}}$, define

$$[\mathbf{K}_j]_{ab} := k_j^{(0)}\big(\widehat{\boldsymbol{\mu}}^{(a)}, \widehat{\boldsymbol{\mu}}^{(b)}\big) + (\sigma_{n,j}^2 + \alpha_j)\delta_{ab}, \qquad [\mathbf{k}_j(\widehat{\boldsymbol{\mu}})]_a := k_j^{(0)}\big(\widehat{\boldsymbol{\mu}}^{(a)}, \widehat{\boldsymbol{\mu}}\big). \tag{5.25}$$

The marginal posterior mean and predictive variance at a standardized query are

$$\widehat{m}_j(\widehat{\boldsymbol{\mu}}) = \mathbf{k}_j(\widehat{\boldsymbol{\mu}})^\top \mathbf{K}_j^{-1} \widehat{\mathbf{z}}_{j,\mathrm{tr}}, \tag{5.26}$$

$$\widehat{v}_j(\widehat{\boldsymbol{\mu}}) = k_j^{(0)}(\widehat{\boldsymbol{\mu}}, \widehat{\boldsymbol{\mu}}) + \sigma_{n,j}^2 - \mathbf{k}_j(\widehat{\boldsymbol{\mu}})^\top \mathbf{K}_j^{-1} \mathbf{k}_j(\widehat{\boldsymbol{\mu}}). \tag{5.27}$$

The inverse notation represents covariance-matrix solves; the implementation does not form an explicit inverse. If the original coordinate standard deviation is $s_{z,j}$, inverse scaling gives

$$m_j(\boldsymbol{\mu}) = \overline{z}_j + s_{z,j} \widehat{m}_j(\widehat{\boldsymbol{\mu}}), \qquad v_j(\boldsymbol{\mu}) = s_{z,j}^2 \widehat{v}_j(\widehat{\boldsymbol{\mu}}).$$

The deterministic POD–GPR displacement used in the numerical comparisons is reconstructed from the posterior coordinate means. Because the coordinate processes are fitted independently, the reported uncertainties are marginal coordinate-wise quantities and contain no cross-coordinate posterior covariance. They are not certified bounds for the finite element field error and do not include POD truncation or broken-to-continuous output-transfer error. The common offline–online workflow and the project-specific implementation settings are summarized in Section A.5.

## 5.4 Assessment protocol and offline–online accounting

The numerical assessment uses a common error and timing convention for all methods. Validation fields are generated by fresh FOM solves at parameter values unseen by the predictive models and lying within the declared admissible parameter domains. They support retained-dimension and model-setting decisions, whereas the reported test evaluations remain separate from this selection step [54]. For a test parameter $\boldsymbol{\mu}$, the relative displacement error is

$$e_{L^2}(\boldsymbol{\mu}) := \frac{\|\widetilde{w}_h(\boldsymbol{\mu}) - \widetilde{w}_N(\boldsymbol{\mu})\|_{L^2(\Omega)}}{\|\widetilde{w}_h(\boldsymbol{\mu})\|_{L^2(\Omega)}}. \tag{5.28}$$

Here both fields are represented in the same global continuous comparison space defined by Equation (5.14); for a panel partition, the panel-wise physical field is assembled before that projection. For the predictive non-intrusive methods, $\widetilde{w}_N$ is reconstructed from the predicted POD coordinates; for POD–Proj it is the direct retained-space projection, while the lifted POD–G field is transferred to the same comparison space. Thus all methods are assessed using one physical $L^2(\Omega)$ output representation.

For each test query $i$, $t_{\mathrm{FOM}}^{(i)}$ is the wall-clock time for parameter assignment, setup of the full-order variational problem, the nonlinear Newton solve, and construction of the returned displacement field, whereas $t_{\mathrm{ROM}}^{(i)}$ measures the online reduced evaluation. For POD–G, it includes parameter assignment and problem setup, full-order residual and Jacobian assembly at each reduced Newton iteration, solution of the reduced linear systems, and lifting of the reduced solution. For PODI–RBF, PODI–Linear, POD–GPR, and POD–NN, it includes componentwise parameter scaling, prediction of the POD coordinates, inverse coordinate scaling, and field

reconstruction. The per-query speedup is

$$S_t^{(i)} := \frac{t_{\mathrm{FOM}}^{(i)}}{t_{\mathrm{ROM}}^{(i)}}. \tag{5.29}$$

Timing begins after the offline models have been constructed and loaded, and excludes error evaluation and plotting. Over a test set of $N_{\mathrm{test}}$ parameters, the reported averages are

$$\overline{S}_t := \frac{1}{N_{\mathrm{test}}} \sum_{i=1}^{N_{\mathrm{test}}} S_t^{(i)}, \qquad \overline{e}_{L^2} := \frac{1}{N_{\mathrm{test}}} \sum_{i=1}^{N_{\mathrm{test}}} e_{L^2}(\boldsymbol{\mu}^{(i)}). \tag{5.30}$$

For POD–Proj, $t_{\mathrm{ROM}}^{(i)}$ comprises only projection and lifting of an available FOM state, so the associated ratio quantifies the cost of the POD-space representation step.

## 5.5 ROM assessment for high-dimensional mechanical parametrization

The first benchmark considers a monolithic plate with a six-dimensional parameter vector. It compares the intrusive POD–Galerkin model with the four non-intrusive surrogates and quantifies the resulting accuracy–cost trade-off for a nonlinear plate problem with simultaneous variation of stiffness and loading. The full-order model is the one defined in Chapter 4, with free exterior boundaries and parameter vector

$$\boldsymbol{\mu} = (D_x, D_y, D_{xy}, D_s, k_s, f) \in \mathbb{R}^6. \tag{5.31}$$

The offline snapshot set contains $N_s = 1000$ Latin-hypercube samples over

$$D_x, D_y \in [5 \times 10^3,\, 2 \times 10^4]\,\mathrm{N\,m}, \quad D_{xy}, D_s \in [10^3,\, 5 \times 10^3]\,\mathrm{N\,m}, \tag{5.32}$$

$$k_s \in [10^6,\, 10^7]\,\mathrm{N/m^3}, \qquad f \in [-10^4,\, -2 \times 10^3]\,\mathrm{N/m^2}. \tag{5.33}$$

The left panel of Figure 5.1 shows the scores of the first two principal components obtained by applying principal component analysis (PCA) to the componentwise-standardized six-dimensional LHS design, with each point representing one offline parameter sample. The first and second principal components explain 16.8% and 16.7% of the standardized parameter variance, respectively. Their combined contribution of 33.5% indicates that the variability of the input design is distributed across several parameter directions rather than concentrated in a dominant one- or two-dimensional linear subspace. The score plane therefore provides a compact visualization of the joint distribution of the offline samples along the two directions of largest variance, while the LHS construction supplies the marginal stratification over each parameter range.

For each offline parameter, the FOM is solved and the corresponding displacement field is stored in the snapshot matrix. The POD diagnostic uses $\eta_{\mathrm{POD}} = 0.999$. The normalized singular values in the right panel of Figure 5.1 decay by several orders of magnitude, and the first five modes contain more than 99.9% of the snapshot energy. Thus, although

the parameter samples vary along several independent directions, the associated displacement responses lie close to a low-dimensional linear subspace, providing the basis for an efficient POD approximation.

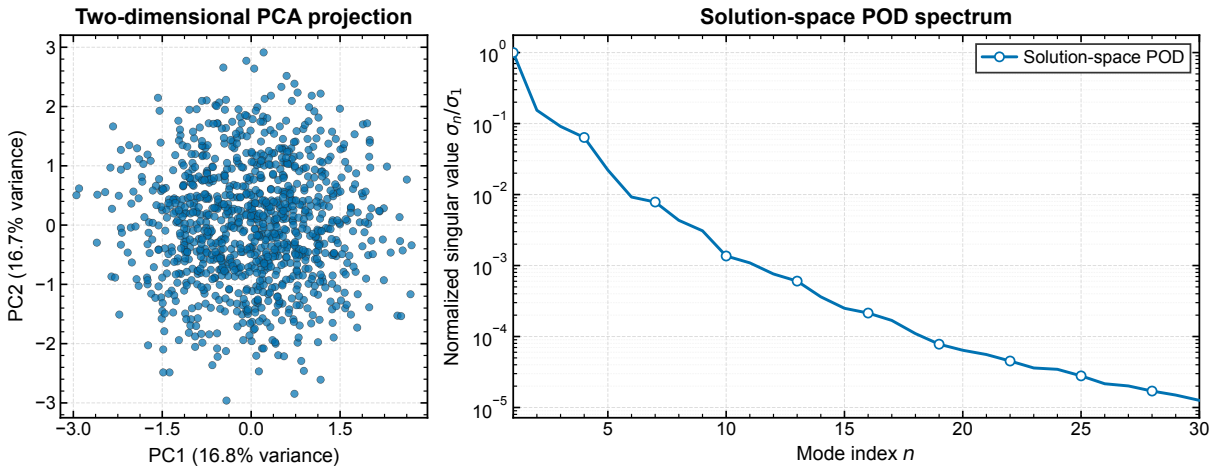



**Figure 5.1:** Offline sampling and POD diagnostics for the high-dimensional mechanical ROM test. Left: first two principal-component scores of the 1000 componentwise-standardized Latin-hypercube samples, accounting jointly for 33.5% of the standardized parameter variance. Right: decay of the normalized solution-space singular values.

For a prescribed dimension $N$, POD–G determines the reduced coordinates from the projected equilibrium equations, whereas PODI–RBF, PODI–Linear, POD–GPR, and POD–NN approximate the projection-coordinate map introduced in Section 5.3. The comparison therefore uses the same number of POD modes while varying the mechanism by which the online coordinates are obtained.

The singular-value decay and retained-energy criterion quantify the compressibility of the snapshot space. To examine the influence of basis enrichment on the complete reduced models, Figure 5.2 reports the validation error for several prescribed dimensions. For each value of $N$, the first $N$ ordered POD modes define the corresponding POD–G and predictive models, which are evaluated on the same validation set.

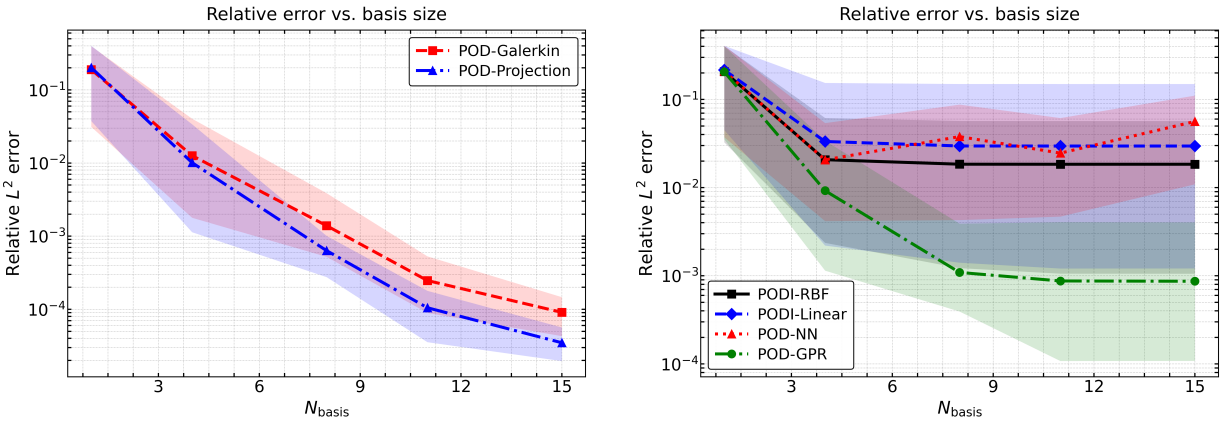



**Figure 5.2:** Validation-error sensitivity to the retained POD dimension for the high-dimensional mechanical ROM test. Left: intrusive POD–Galerkin error and POD projection reference. Right: non-intrusive predictive ROM errors. Curves show the mean relative $L^2(\Omega)$ error, and the shaded bands span the minimum and maximum errors over the validation samples.

For the intrusive comparison in Figure 5.2, POD–Proj gives the representation error obtained by projecting the FOM field onto the retained basis, whereas POD–G determines the coordinates from the reduced equilibrium equations. Because the two curves use the same POD space, their

separation measures the additional error associated with the reduced equilibrium solve. The mean POD–G error decreases from the $10^{-1}$ range at $N = 1$ to the $10^{-3}$ range by $N = 8$, and falls below $10^{-4}$ at the largest tested dimension.

The non-intrusive behavior is more heterogeneous. POD–GPR benefits most from enrichment and reaches the $10^{-3}$ range. PODI–RBF improves markedly up to approximately $N = 4$ and then levels off in the $10^{-2}$ range. PODI–Linear and POD–NN gain comparatively little from the richest bases. Once the POD-space representation error is small, interpolation or regression error dominates these predictive approximations.

The singular-value decay and validation trends motivate $N = 4$ and $N = 8$ as representative compact and enriched configurations, respectively. The final comparison evaluates both configurations on the same $N_{\text{test}} = 100$ test parameters. Figure 5.3 presents the sample-wise relative errors, with $N = 4$ on the left and $N = 8$ on the right. Within each panel, the samples are ranked by $\|\widehat{\boldsymbol{\mu}}^{(i)}\|_2$, where $\widehat{\boldsymbol{\mu}}^{(i)}$ is the componentwise-standardized test parameter. This dimensionless rank supplies a reproducible display order, while Table 5.1 reports the mean errors and speedups.

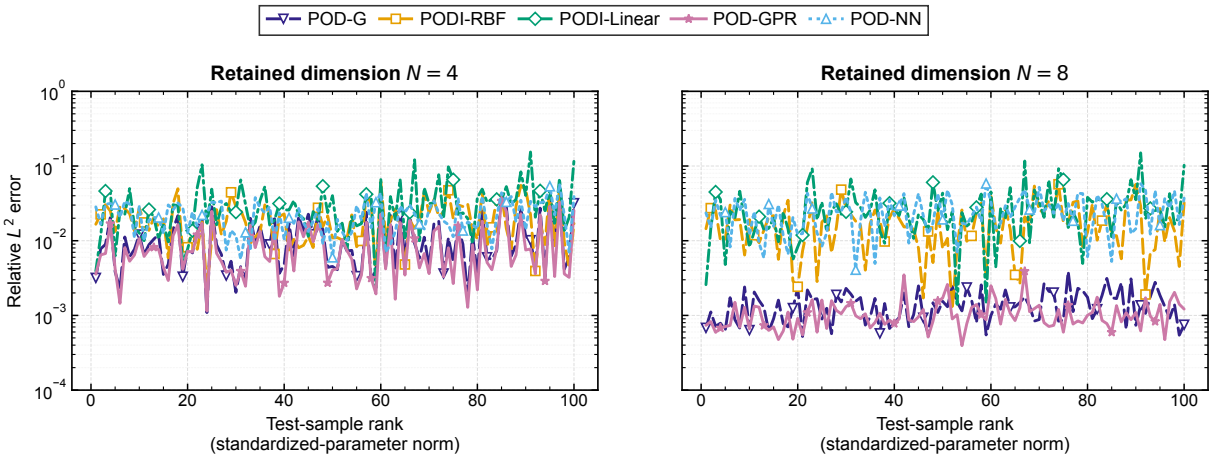



**Figure 5.3:** Sample-wise relative $L^2(\Omega)$ errors over $N_{\text{test}} = 100$ test parameters. Left: retained dimension $N = 4$. Right: retained dimension $N = 8$. In both panels, the sample rank follows the Euclidean norm of the componentwise-standardized test parameter vector.

**Table 5.1:** Mean performance for the high-dimensional mechanical ROM test over $N_{\text{test}} = 100$ test parameters for the retained dimensions $N = 4$ and $N = 8$. The mean speedup $\overline{S}_t$ is the arithmetic mean of the per-query ratios defined in Equation (5.29).

| **Method** | **Basis size** $N$ | **Mean** $e_{L^2}$ | **Mean speedup** $\overline{S}_t$ |
|---|---|---|---|
| POD–G | 4 | $1.15 \times 10^{-2}$ | 4.10 |
| | 8 | $1.43 \times 10^{-3}$ | 3.69 |
| PODI–RBF | 4 | $2.07 \times 10^{-2}$ | $3.73 \times 10^{3}$ |
| | 8 | $1.84 \times 10^{-2}$ | $3.51 \times 10^{3}$ |
| PODI–Linear | 4 | $3.34 \times 10^{-2}$ | $3.82 \times 10^{3}$ |
| | 8 | $2.97 \times 10^{-2}$ | $3.59 \times 10^{3}$ |
| POD–GPR | 4 | $9.23 \times 10^{-3}$ | $1.46 \times 10^{3}$ |
| | 8 | $1.09 \times 10^{-3}$ | $8.96 \times 10^{2}$ |
| POD–NN | 4 | $2.28 \times 10^{-2}$ | $1.90 \times 10^{3}$ |
| | 8 | $2.87 \times 10^{-2}$ | $1.75 \times 10^{3}$ |

At $N = 4$, POD–GPR has the smallest predictive mean error, $9.23 \times 10^{-3}$, while POD–G gives $1.15 \times 10^{-2}$. At $N = 8$, their means decrease to $1.09 \times 10^{-3}$ and $1.43 \times 10^{-3}$, respectively. PODI–RBF and PODI–Linear improve

moderately with enrichment, while POD–NN changes marginally; their errors are therefore governed mainly by the coordinate surrogate.

All predictive non-intrusive methods provide substantial mean speedup. PODI–RBF and PODI–Linear are the fastest, with factors of about $3.5\times10^3$–$3.8\times10^3$. POD–NN reaches approximately $1.7\times10^3$–$1.9\times10^3$, and POD–GPR approximately $9.0\times10^2$–$1.5\times10^3$. POD–G gives a factor of about four because its online stage assembles the full-order residual and Jacobian.

Figure 5.4 places the FOM displacement alongside the absolute-error fields of POD–G, PODI–RBF, PODI–Linear, POD–GPR, and POD–NN for a representative test parameter. Each field has its own color scale, allowing the spatial structure of the method-specific errors to remain visible across their different magnitudes.

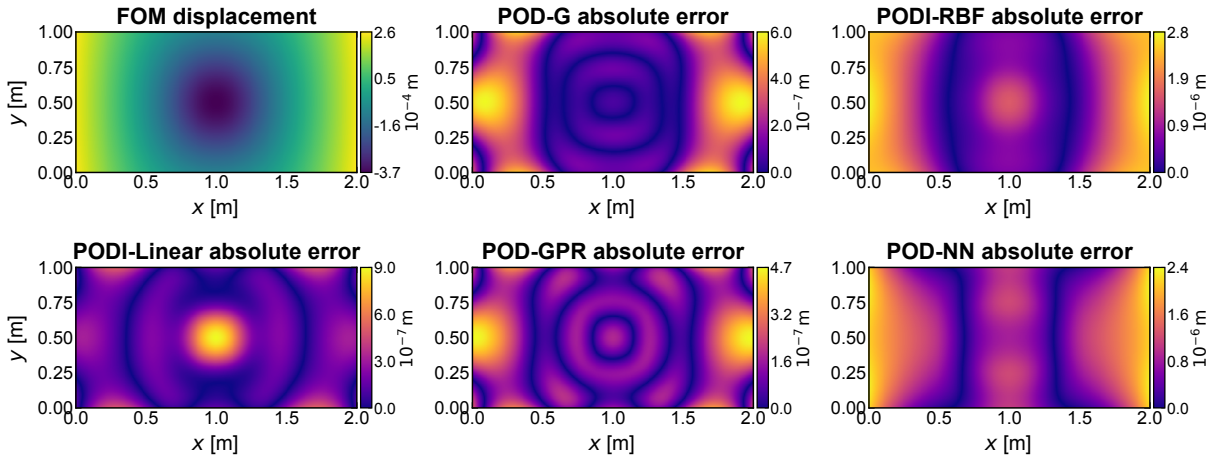



**Figure 5.4:** Representative field comparison for the high-dimensional mechanical ROM test. Top row: FOM displacement, POD–G absolute error, and PODI–RBF absolute error. Bottom row: PODI–Linear, POD–GPR, and POD–NN absolute errors. The FOM field and each absolute-error field use an independent color scale, indicated by the adjacent colorbar; quantitative comparisons follow the corresponding colorbar values.

For this representative parameter, the maximum absolute errors range from approximately $4.7\times10^{-7}$ m to $2.8\times10^{-6}$ m. The independently normalized fields resolve the method-dependent spatial patterns, while the complete test-set statistics in Table 5.1 quantify the overall predictive performance.

## 5.6 ROM assessment for penalty-coupled panel systems

The second benchmark considers non-intrusive surrogates for a two-panel partition with one physical panel interface. The structure is a rectangular plate divided into two side-by-side panels, $(\Omega_1, \Omega_2)$, separated by a vertical interface. All exterior boundaries are free, a localized load patch acts on the plate, and displacement and normal-rotation continuity are imposed weakly by the interface terms of the full-order formulation. The parameter vector is

$$\boldsymbol{\mu} = (k_s, f) \in \mathbb{R}^2, \tag{5.34}$$

with $k_s \in [10^6, 10^7]\ \mathrm{N/m^3}$ and $f \in [-10^4, -2\times10^3]\ \mathrm{N/m^2}$.

An offline Latin-hypercube design with $N_s = 300$ samples is used. The left panel of Figure 5.5 shows the principal-component scores of the componentwise-standardized parameter pairs. Since $\boldsymbol{\mu} = (k_s, f) \in \mathbb{R}^2$,

PCA yields two orthogonal principal directions that span the complete centered standardized parameter variation. PC1 and PC2 account for 50.1% and 49.9% of the total standardized parameter variance, respectively, and therefore account jointly for all variance of this two-dimensional standardized parameter dataset. The approximately diamond-shaped score cloud results from expressing the standardized rectangular Latin-hypercube design in the rotated orthogonal coordinate system defined by PCA.

For each offline parameter, the two-panel FOM is solved and the transverse displacement is transferred to the fixed continuous output space through Equation (5.14). With $\eta_{\mathrm{POD}} = 0.999$, the normalized projected-space singular values in the right panel of Figure 5.5 decay rapidly and reach approximately $10^{-6}$ by the eighth mode.

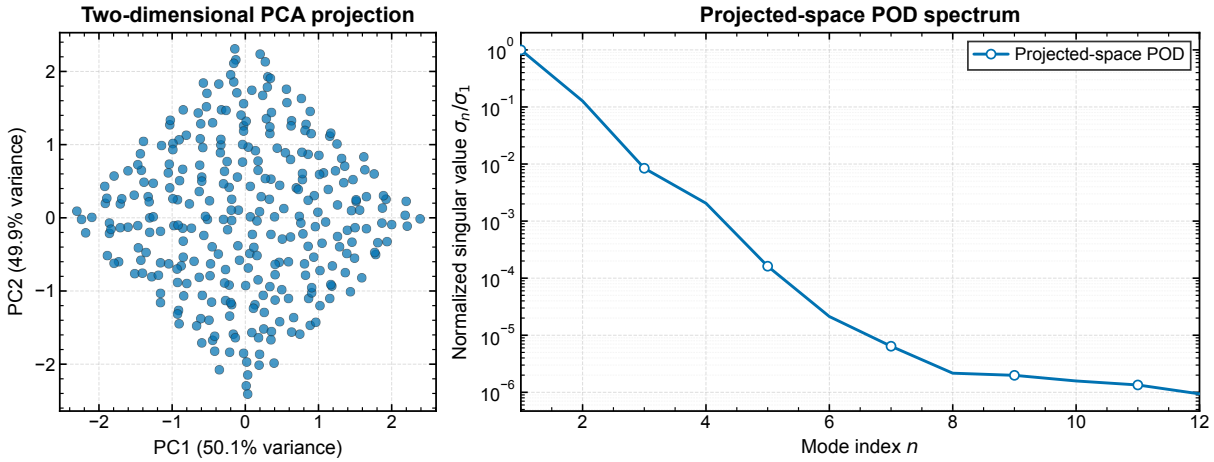



**Figure 5.5:** Offline sampling and POD diagnostics for the two-panel mechanical ROM test. Left: PCA scores of the $N_s = 300$ componentwise-standardized Latin-hypercube samples; PC1 and PC2 account for 50.1% and 49.9% of the standardized parameter variance, respectively, and the diamond-shaped cloud reflects the PCA rotation of the standardized rectangular sampling domain. Right: decay of the normalized projected-space POD singular values.

The selected two-panel configuration uses a transferred-space POD basis of dimension $N = 8$. To illustrate the effect of basis enrichment, the representative profile comparison also includes the compact configuration $N = 1$. The predictive comparison comprises PODI–RBF, PODI–Linear, POD–GPR, and POD–NN, while POD–Proj provides the POD-space projection reference.

Figure 5.6 compares transverse-deflection profiles along a mid-span cross-section and the plate diagonal for the FOM, POD–Proj, and the four predictive ROMs at the same representative test parameter, using the compact $N = 1$ and selected $N = 8$ configurations. At $N = 1$, the reduced models reproduce the overall deformation shape, although small deviations from the FOM remain visible along both profiles. Enriching the basis to $N = 8$ markedly reduces these discrepancies, with the POD–Proj and predictive ROM profiles lying in close agreement with the FOM. The comparison therefore illustrates the improvement in the reduced representation obtained by increasing the retained POD dimension.

The corresponding field comparison in Figure 5.7 resolves the differences among the reduced representations at the selected dimension $N = 8$. PODI–RBF and POD–GPR have maximum absolute errors of approximately $1.6 \times 10^{-8}$ m and $5.8 \times 10^{-9}$ m, respectively, whereas PODI–Linear and POD–NN reach approximately $1.6 \times 10^{-6}$ m and $4.2 \times 10^{-6}$ m. The FOM displacement provides the common physical reference, and the

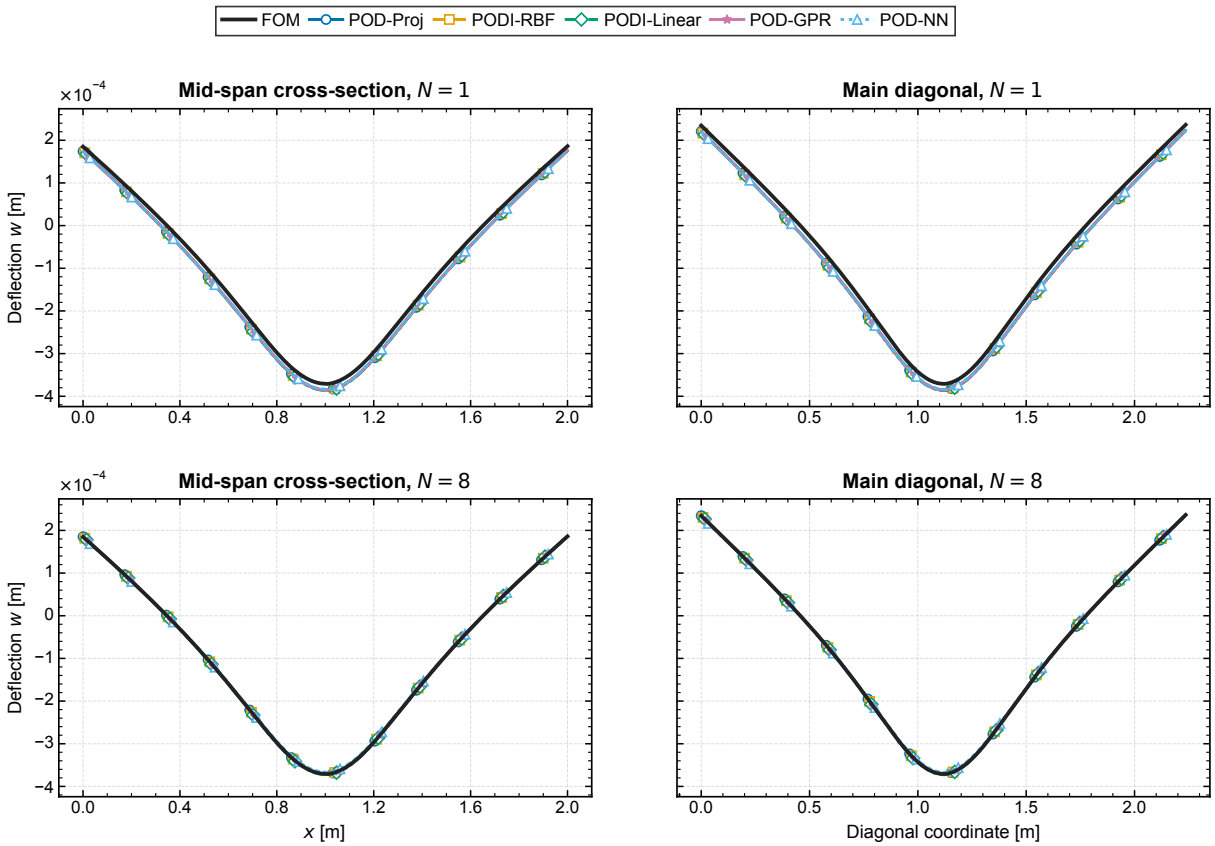



**Figure 5.6:** Representative transverse-deflection profiles for the FOM, POD–Proj, and the four predictive ROMs at retained dimensions $N = 1$ and $N = 8$. Top row: $N = 1$; bottom row: $N = 8$. Left: mid-span cross-section. Right: plate diagonal. The same representative parameter is used for both retained dimensions.

independently scaled error fields make the method-dependent spatial structures visible at their respective magnitudes.

For the quantitative assessment, the models are evaluated on $N_{\text{test}} = 100$ parameters drawn from the same ranges as the offline set. Figure 5.8 reports the sample-wise relative errors, with samples ranked by the norm of the componentwise-standardized test parameter vector. The corresponding mean speedups are listed in Table 5.2.

The POD–Proj mean error, $4.24\times10^{-6}$, gives the POD-space representation reference. Among the predictive surrogates, POD–GPR and PODI–RBF give the two smallest mean errors, $8.33\times10^{-6}$ and $1.86\times10^{-5}$, respectively. PODI–Linear and POD–NN yield larger means of order $10^{-3}$ and $10^{-2}$. PODI–RBF and PODI–Linear are the fastest predictive methods, with mean speedups near $2.1 \times 10^4$, while POD–GPR and POD–NN provide factors of $7.66 \times 10^3$ and $1.01 \times 10^4$.

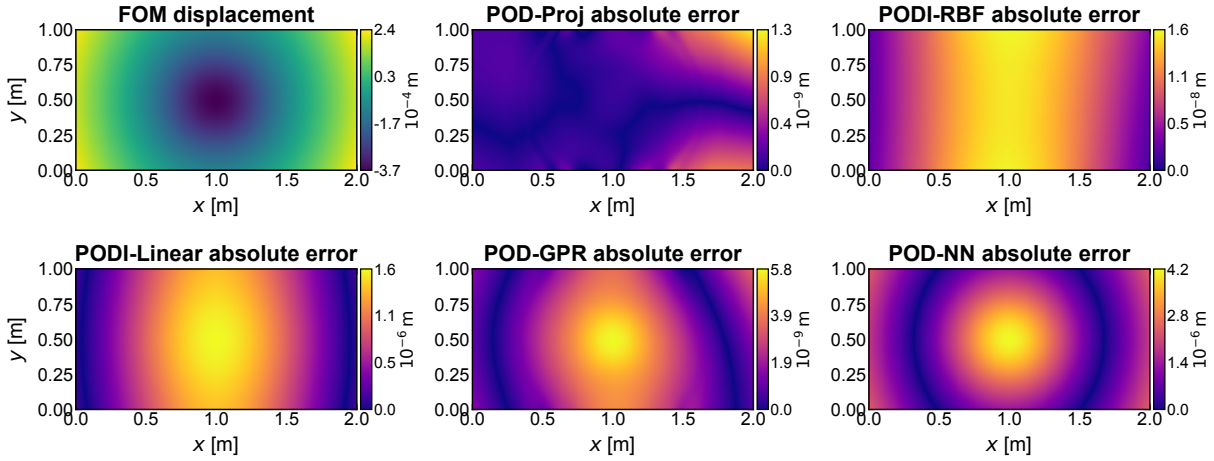



**Figure 5.7:** Representative field comparison for the two-panel mechanical ROM test at the selected retained dimension $N = 8$. Top row: FOM displacement, POD–Proj absolute error, and PODI–RBF absolute error. Bottom row: PODI–Linear, POD–GPR, and POD–NN absolute errors. Displacement and absolute errors are reported in meters, with each field using its displayed color scale.

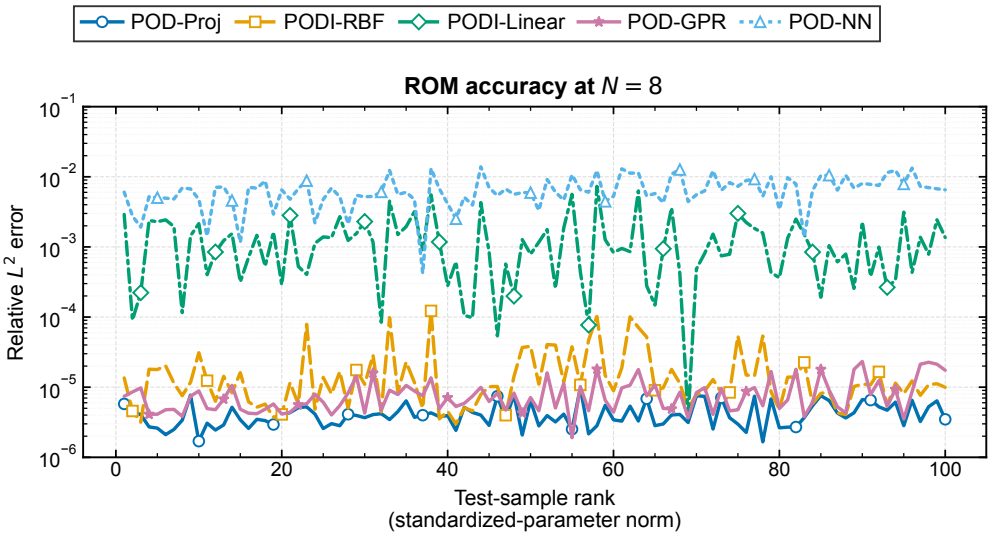



**Figure 5.8:** Sample-wise relative $L^2(\Omega)$ errors for the two-panel system at $N = 8$ over $N_{\text{test}} = 100$ parameters, ranked by the norm of the componentwise-standardized test parameter vector. POD–Proj provides the POD-space projection reference.

**Table 5.2:** Mean performance for the two-panel penalty-coupled system over $N_{\text{test}} = 100$ test parameters at the selected retained dimension $N = 8$. The POD–Proj timing corresponds to projection and lifting from an available FOM state.

| **Method** | **Mean** $e_{L^2}$ | **Mean speedup** $\overline{S}_t$ |
|---|---|---|
| POD–Proj (diagnostic) | $4.24 \times 10^{-6}$ | $1.97 \times 10^{3}$ |
| PODI–RBF | $1.86 \times 10^{-5}$ | $2.08 \times 10^{4}$ |
| PODI–Linear | $1.40 \times 10^{-3}$ | $2.06 \times 10^{4}$ |
| POD–GPR | $8.33 \times 10^{-6}$ | $7.66 \times 10^{3}$ |
| POD–NN | $6.78 \times 10^{-3}$ | $1.01 \times 10^{4}$ |

## 5.7 ROM assessment for multi-panel configurations

The final benchmark evaluates the non-intrusive framework on two penalty-coupled multi-panel configurations with different boundary conditions and load patterns. Both use $\boldsymbol{\mu} = (k_s, f) \in \mathbb{R}^2$, with $k_s \in [10^6, 10^7]\ \mathrm{N/m^3}$ and $f \in [-10^4, -2 \times 10^3]\ \mathrm{N/m^2}$. For each configuration, an offline snapshot set of $N_s = 300$ parameters is generated by Latin hypercube sampling, and the corresponding FOM solutions are transferred to the fixed continuous output space through Equation (5.14).

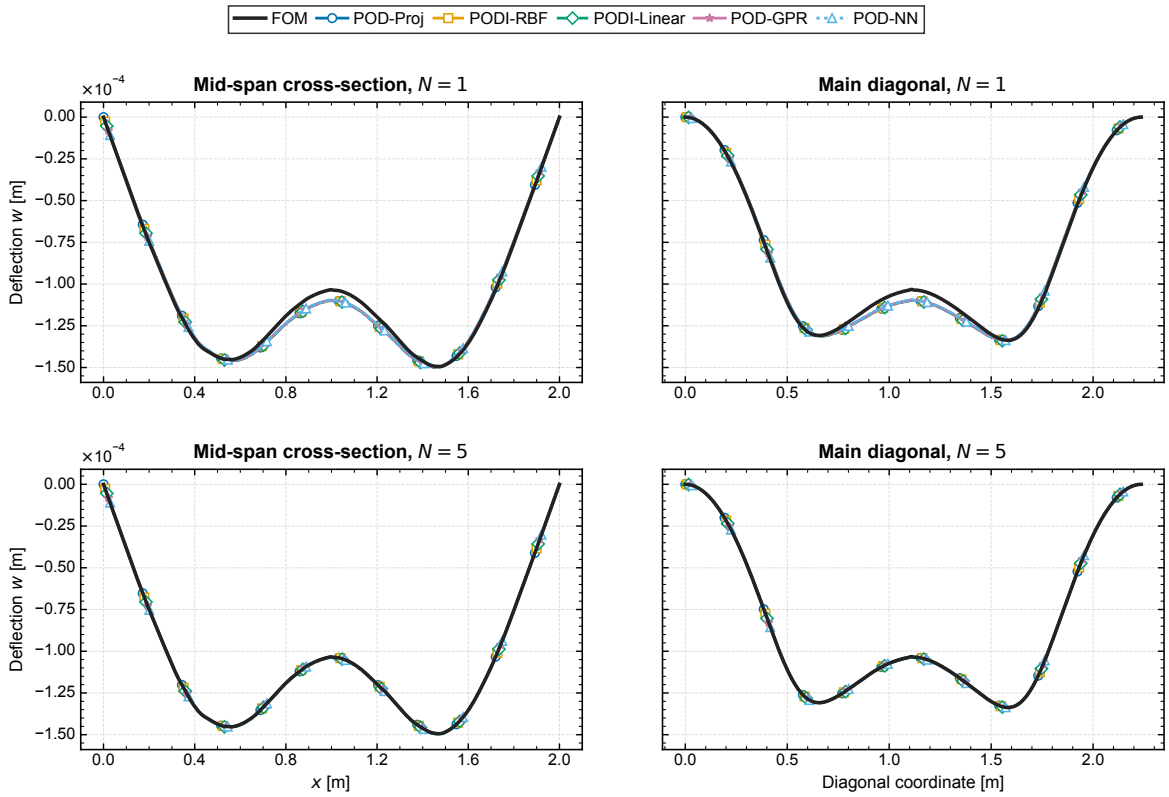
The first configuration partitions the plate into four panels, $\Omega = \bigcup_{i=1}^{4} \Omega_i$, in a $2 \times 2$ layout. The exterior edges are simply supported, and multiple localized load patches produce strong curvature variations near the physical panel interfaces. With $\eta_{\mathrm{POD}} = 0.999$, the retained dimension is $N = 5$, and these modes contain more than 99.99% of the snapshot energy. The second configuration consists of three side-by-side panels, $\Omega = \bigcup_{i=1}^{3} \Omega_i$, separated by vertical interfaces. Its exterior edges are free, the transverse load varies linearly in the $x$-direction, and the Winkler foundation provides global support. The same POD threshold gives $N = 3$.

The predictive models for both configurations are PODI–RBF, PODI–Linear, POD–GPR, and POD–NN, all of which approximate the projection-coordinate map introduced in Section 5.3. To assess the influence of basis enrichment consistently with the two-panel benchmark, the representative profile comparisons additionally consider the compact retained dimension $N = 1$, alongside the selected dimensions $N = 5$ for the four-panel system and $N = 3$ for the three-panel system. POD–Proj supplies the POD-space projection reference.

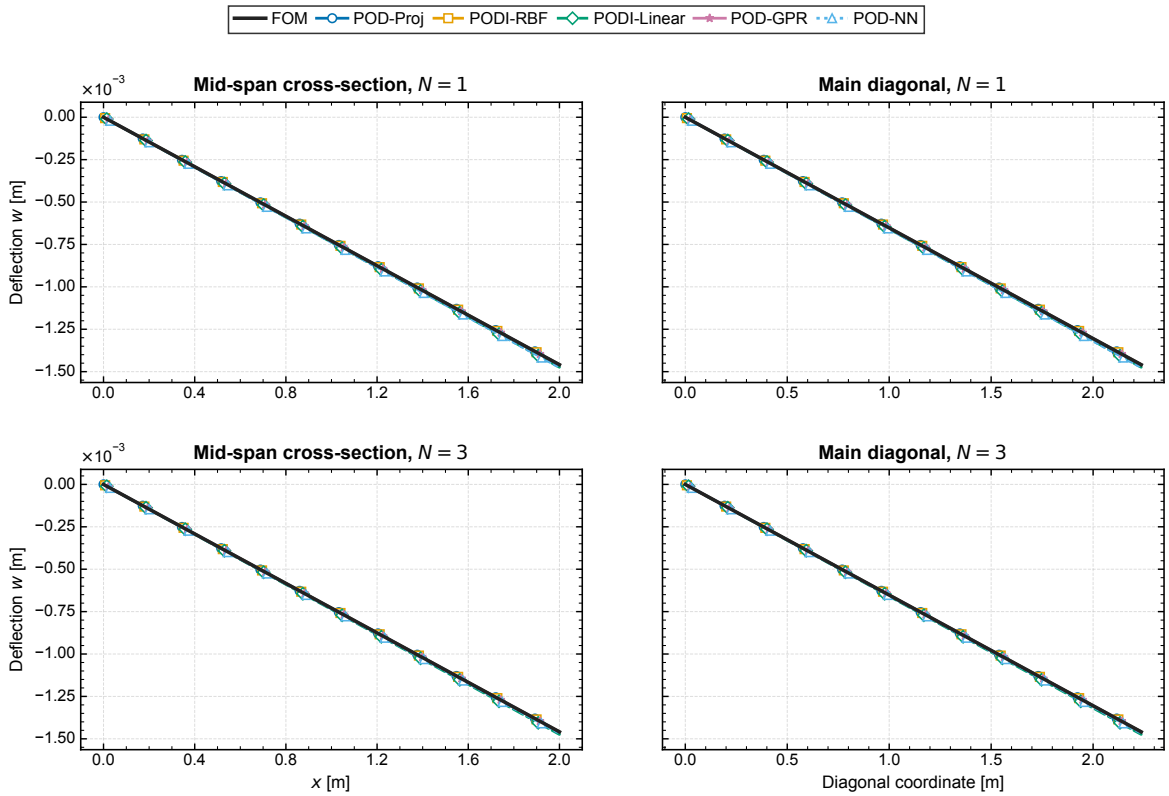
Figure 5.9 compares the corresponding transverse-deflection profiles along a central cross-section and the plate diagonal. For the four-panel simply supported configuration, $N = 1$ already reproduces the global deformation pattern, while small deviations from the FOM remain visible in regions of stronger curvature. Enrichment to $N = 5$ reduces these discrepancies and brings the POD–Proj and predictive ROM profiles into close agreement with the FOM. For the three-panel free-edge configuration, the $N = 1$ profiles are already close to the FOM at the plotted scale, and increasing the retained dimension to $N = 3$ produces only a modest further change. The two configurations therefore differ in their sensitivity to enrichment of the retained POD space.

Figure 5.10 complements the profile comparison by showing the FOM displacement and the spatial absolute-error fields at the selected retained dimension for each configuration. For the four-panel system at $N = 5$, the maximum absolute errors are approximately $9.6 \times 10^{-12}$ m for POD–Proj, $3.2 \times 10^{-9}$ m for PODI–RBF, $3.2 \times 10^{-7}$ m for PODI–Linear, $1.2 \times 10^{-9}$ m for POD–GPR, and $1.2 \times 10^{-7}$ m for POD–NN. For the three-panel system at $N = 3$, the corresponding maxima are approximately $7.5 \times 10^{-10}$ m, $1.9 \times 10^{-7}$ m, $1.5 \times 10^{-5}$ m, $7.9 \times 10^{-9}$ m, and $8.0 \times 10^{-6}$ m, respectively. The independently scaled fields make the different spatial error patterns visible despite the several orders of magnitude separating their amplitudes.

For each configuration, the models are evaluated on $N_{\mathrm{test}} = 100$ test parameters sampled from the same ranges as the offline set. Figure 5.11 reports the sample-wise relative errors for the four-panel and three-panel systems, using the standardized-parameter-norm ranking introduced




*Four-panel simply supported system: $N = 1$ and selected $N = 5$.*

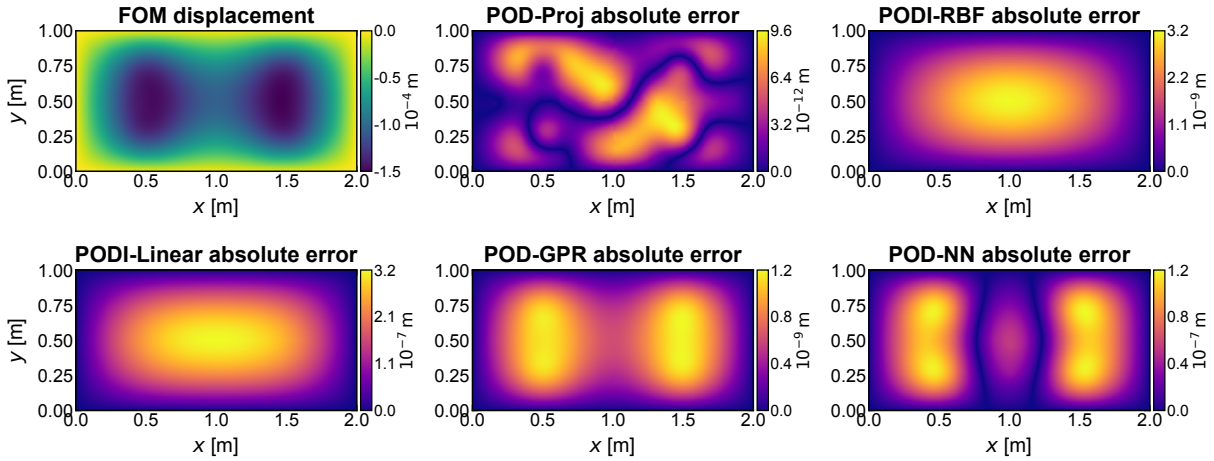



*Three-panel free-edge system: $N = 1$ and selected $N = 3$.*

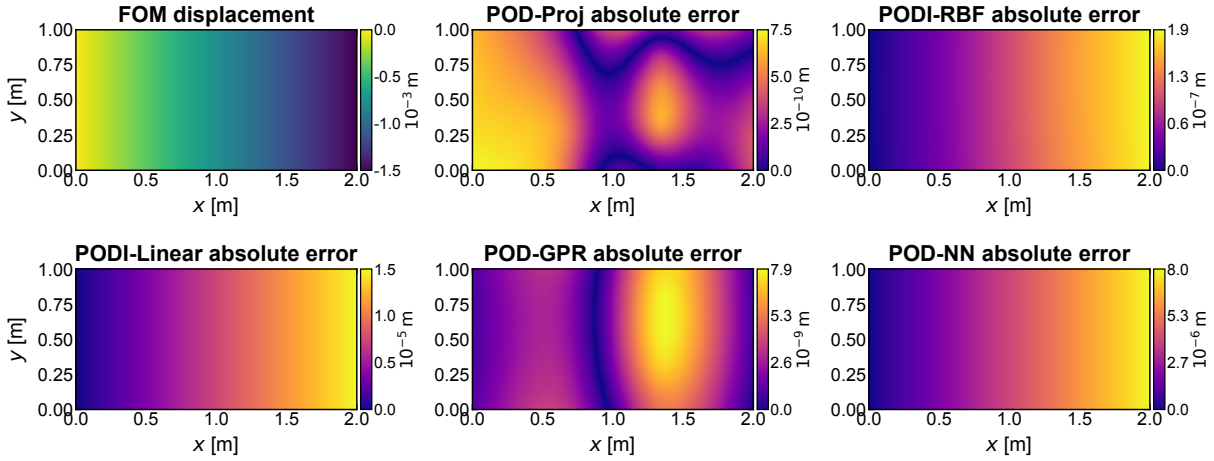
**Figure 5.9:** Representative transverse-deflection profiles for the FOM, POD–Proj, and the four predictive ROMs. In each configuration, the top and bottom rows show $N = 1$ and the selected retained dimension, respectively, while the left and right panels show the central cross-section and plate diagonal for the same representative parameter.

in Section 5.5. The corresponding mean speedups are reported in Table 5.3.

Across both multi-panel systems, POD–GPR has the smallest predictive mean error, $4.81 \times 10^{-6}$ and $7.86 \times 10^{-6}$, followed by PODI–RBF with $8.38 \times 10^{-6}$ and $3.76 \times 10^{-5}$. PODI–Linear and POD–NN have larger mean errors of order $10^{-3}$–$10^{-2}$. All four predictive surrogates attain mean speedups between $2.35 \times 10^{4}$ and $5.03 \times 10^{4}$, with PODI–RBF and PODI–Linear giving the largest or nearly largest factors.




*Four-panel simply supported system at the selected retained dimension $N = 5$.*




*Three-panel free-edge system at the selected retained dimension $N = 3$.*

**Figure 5.10:** Representative field comparisons for the two multi-panel mechanical ROM tests at the selected retained dimensions. Within each configuration, the top row shows the FOM displacement, POD–Proj absolute error, and PODI–RBF absolute error; the bottom row shows the PODI–Linear, POD–GPR, and POD–NN absolute errors. Displacement and absolute errors are reported in meters, with each field using its displayed color scale.

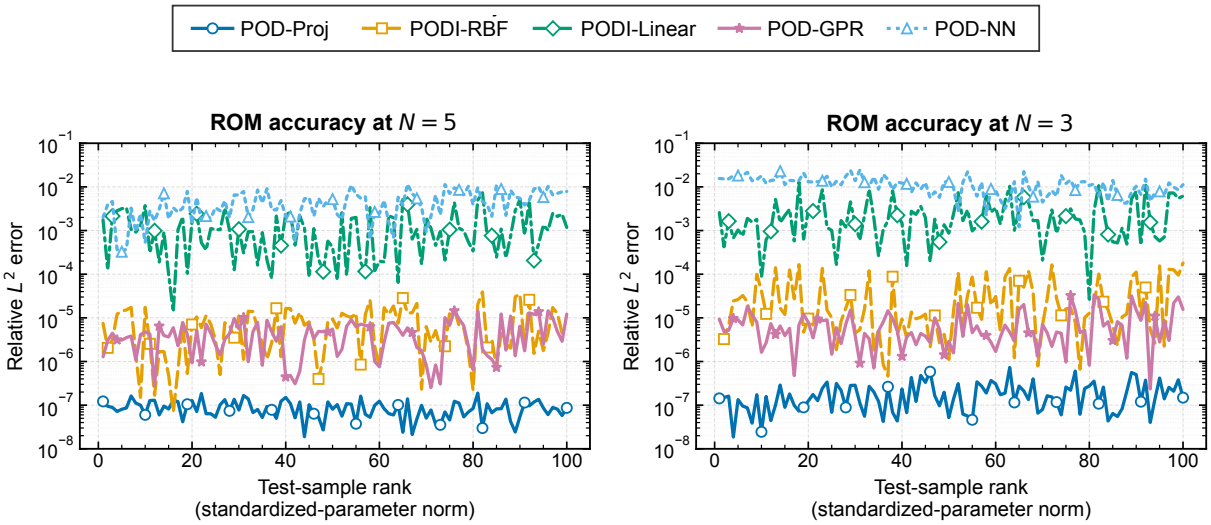



**Figure 5.11:** Sample-wise relative $L^2(\Omega)$ errors over $N_{\text{test}} = 100$ parameters for the two multi-panel systems. Left: four-panel simply supported system ($N = 5$). Right: three-panel free-edge system ($N = 3$). Samples are ranked by the norm of the componentwise-standardized test parameter vector; POD–Proj provides the POD-space projection reference. A single legend applies to both panels.

**Table 5.3:** Mean performance for the two multi-panel systems over $N_{\text{test}} = 100$ test parameters per configuration. The POD–Proj timing corresponds to projection and lifting from an available FOM state.

| | **Four-panel simply supported** $N = 5$ | | **Three-panel free-edge** $N = 3$ | |
|---|---|---|---|---|
| **Method** | **Mean** $e_{L^2}$ | **Mean speedup** $\overline{S}_t$ | **Mean** $e_{L^2}$ | **Mean speedup** $\overline{S}_t$ |
| POD–Proj (diagn.) | $8.75 \times 10^{-8}$ | $2.46 \times 10^{3}$ | $1.88 \times 10^{-7}$ | $2.44 \times 10^{3}$ |
| PODI–RBF | $8.38 \times 10^{-6}$ | $5.01 \times 10^{4}$ | $3.76 \times 10^{-5}$ | $4.75 \times 10^{4}$ |
| PODI–Linear | $1.37 \times 10^{-3}$ | $5.03 \times 10^{4}$ | $2.69 \times 10^{-3}$ | $4.75 \times 10^{4}$ |
| POD–GPR | $4.81 \times 10^{-6}$ | $2.36 \times 10^{4}$ | $7.86 \times 10^{-6}$ | $2.74 \times 10^{4}$ |
| POD–NN | $4.70 \times 10^{-3}$ | $2.49 \times 10^{4}$ | $1.17 \times 10^{-2}$ | $2.35 \times 10^{4}$ |

## 5.8 Summary and limitations of the mechanical ROMs

The mechanical ROM study constructs and assesses reduced-order models for the high-fidelity problem developed in Chapters 3 and 4. Metric-aware POD is applied in the native full-order space for intrusive reduction. For non-intrusive learning and common field comparison, panel-wise displacement is transferred by an explicit $L^2$-projection to a fixed global continuous space, and the physical parameters are mapped to the resulting POD coordinates. The intrusive POD–Galerkin model preserves the residual structure of the high-fidelity problem, whereas PODI–RBF, PODI–Linear, POD–GPR, and POD–NN replace the online nonlinear solve by interpolation or regression of the coordinate map.

The displacement fields generated by the plate–foundation model have compact POD representations across all tested monolithic and panel-partitioned configurations. In the six-dimensional monolithic test, POD–G and POD–GPR give the smallest mean errors: at $N = 8$, POD–GPR gives $1.09 \times 10^{-3}$ compared with $1.43 \times 10^{-3}$ for POD–G, while its mean speedup is $8.96 \times 10^{2}$ compared with 3.69 for the intrusive model. In the two-, three-, and four-panel systems, POD–GPR gives the smallest predictive mean errors and PODI–RBF the second-smallest, while PODI–RBF and

PODI–Linear provide the largest or nearly largest mean speedups.

Rapid singular-value decay establishes state compressibility, while POD–Proj quantifies the retained-space projection floor for an available FOM state. The excess error of POD–G above that floor reflects the reduced equilibrium solve, and the excess error of the predictive surrogates additionally contains the parameter-to-coordinate-map error.

The intrusive ROM retains full-order residual and Jacobian assembly in its online stage, which limits its acceleration. The predictive models depend on the sampled training domain and the quality of the parameter-to-coordinate map, while the common broken-to-continuous transfer fixes the output ordering but does not isolate seam jumps as separate outputs. POD–Proj remains a compression diagnostic based on an available FOM state, and the Gaussian-process posterior variance remains a coordinate-level model-uncertainty indicator. Within this scope, the chapter establishes the mechanical high-fidelity-to-reduced-order workflow that supports the coupled and time-dependent reductions developed later in the thesis.

Part II

# Thermomechanics and Transient Dynamics of Orthotropic Plates

> This part extends the common plate framework in two distinct physical directions. The thermomechanical study couples three-dimensional heat transfer, the plate thermal driver $\theta$, and displacement through a partitioned FOM, then reduces displacement and $\theta$ in separate one-field workflows. The transient-dynamics study adds transverse inertia and transient loading, verifies the active-degree-of-freedom Newmark formulation, and constructs time-conditioned displacement surrogates from disjoint complete trajectories. The two study sequences retain the common plate and panel-interface conventions established earlier while using distinct physical assumptions and reduced outputs.

# Thermomechanical Kirchhoff–Love Plate Formulation

# 6


6.1 Thermomechanical extension of the mechanical model . . . . . 62
6.2 Geometry, panel partition, and three-dimensional thermal domain . . . . . 62
6.3 Thermal subproblem . . . 64
6.4 Weak formulation of the thermal problem . . . . . 65
6.5 Thermal-to-plate reduction 66
6.6 Thermal bending moments . . . . . 66
6.7 Mechanical equilibrium with thermal loading . . . 68
6.8 Partitioned staggered thermomechanical coupling . 70
6.9 Parameter dependence and feedback structure . . . . . 72
6.10 Summary of the thermomechanical formulation . 72


## 6.1 Thermomechanical extension of the mechanical model

The mechanical chapters in Part I formulate, discretize, and reduce the static response of parameterized orthotropic Kirchhoff–Love plates on unilateral Winkler-type foundations. The present chapter extends that framework to thermomechanical bending. Its purpose is to identify the three-dimensional thermal problem, derive the plate-level thermal quantity that drives bending, and define the feedback between lower-face thermal exchange and transverse displacement.

The mechanical backbone remains the model of Chapter 3: the transverse displacement is positive upward, the bending law is aligned orthotropic, the Winkler support distinguishes compression from uplift, and displacement and normal rotation are transmitted weakly across physical panel interfaces. The additional field is the steady three-dimensional temperature in the plate body. Its centered first moment through the thickness produces a scalar plate thermal driver, which enters the mechanical equilibrium through orthotropic thermal bending moments.

The primary formulation is a two-way partitioned coupling [21, 85]. The current displacement determines the contact–gap heat-transfer coefficient on the substrate-facing surface; the resulting temperature field determines the thermal driver; and the thermal driver modifies the plate equilibrium. The thermal solve, thermal-to-plate reduction, mechanical solve, displacement relaxation, and field-update test are therefore iterated to convergence. The finite element realization and solver assessment are developed in Chapter 7.

The thermomechanical formulation is organized around the coupled maps linking the three-dimensional temperature field, the plate thermal driver, and the transverse displacement. Section 6.2 first fixes the midsurface, panel partition, interface notation, and three-dimensional thermal domain. Sections 6.3 and 6.4 then state the strong and weak thermal problems, respectively. Section 6.5 introduces the centered thermal-to-plate reduction, while Sections 6.6 and 6.7 formulate the thermal bending moments and the resulting thermally loaded plate equilibrium, including the physical moment and shear resultants. Finally, Section 6.8 closes the displacement–temperature feedback relation through the partitioned iteration.

## 6.2 Geometry, panel partition, and three-dimensional thermal domain

Let $(\mathbf{e}_x, \mathbf{e}_y, \mathbf{e}_z)$ be an orthonormal basis of $\mathbb{R}^3$, and let $(x, y, z)$ be the associated Cartesian coordinates. The plate midsurface occupies a bounded open set $\Omega \subset \mathbb{R}^2$ with Lipschitz boundary $\partial\Omega$. The plate has thickness

$t > 0$, and the corresponding three-dimensional reference domain is

$$\Omega^t := \Omega \times \left(-\frac{t}{2}, \frac{t}{2}\right) \subset \mathbb{R}^3. \tag{6.1}$$

The coordinate $z$ is centered at the midsurface, so that $z = 0$ corresponds to the reference midsurface. The outward unit normal to $\partial\Omega^t$ is denoted by $\mathbf{N}$.

The non-overlapping panel partition is the same geometric device used in Chapter 3. Thus, $\Omega$ is represented by open subdomains $\{\Omega_i\}_{i=1}^M$ satisfying

$$\overline{\Omega} = \bigcup_{i=1}^{M} \overline{\Omega_i}, \qquad \Omega_i \cap \Omega_j = \emptyset \quad \text{for } i \neq j. \tag{6.2}$$

For adjacent panels, the common interface is

$$\Gamma_{ij} := \partial\Omega_i \cap \partial\Omega_j,$$

and the internal interface skeleton is

$$\Gamma := \bigcup_{1 \le i < j \le M} \Gamma_{ij}. \tag{6.3}$$

On $\Gamma_{ij}$, the outward unit normals to $\Omega_i$ and $\Omega_j$ are denoted by $\mathbf{n}_i$ and $\mathbf{n}_j$, with $\mathbf{n}_j = -\mathbf{n}_i$. For defining jumps, a reference normal $\mathbf{n} = \mathbf{n}_i = -\mathbf{n}_j$ is fixed locally on each interface. This two-dimensional normal is distinct from the three-dimensional outward normal $\mathbf{N}$ on $\partial\Omega^t$.

For a panel-wise scalar field $v = \{v_i\}_{i=1}^M$, possibly discontinuous across interfaces, the jumps on $\Gamma_{ij}$ are

$$[\![v]\!]_{ij} := v_i - v_j, \qquad [\![\partial_{\mathbf{n}} v]\!]_{ij} := \partial_{\mathbf{n}} v_i - \partial_{\mathbf{n}} v_j, \qquad \partial_{\mathbf{n}} v_i := \mathbf{n} \cdot \nabla v_i. \tag{6.4}$$

For a generic scalar field, the second jump in (6.4) is a normal-derivative jump. When $v = w$, where $w$ is the Kirchhoff–Love transverse displacement, it is the jump of the normal rotation across the panel interface under the convention adopted in Chapter 3.

The exterior boundary of the midsurface is decomposed as

$$\partial\Omega = \Gamma_S \cup \Gamma_F, \qquad \Gamma_S \cap \Gamma_F = \emptyset, \tag{6.5}$$

where $\Gamma_S$ denotes the simply supported portion and $\Gamma_F$ the free portion. On smooth portions of $\partial\Omega$, $\mathbf{n} = (n_x, n_y)$ denotes the outward unit normal, $\mathbf{t} = \mathbf{n}^\perp = (-n_y, n_x)$ the unit tangent, and $\partial_s := \mathbf{t} \cdot \nabla$ the tangential derivative.

For the thermal problem, the boundary of $\Omega^t$ is split into the exposed top face, the substrate-facing bottom face, and the lateral boundary:

$$\begin{aligned} \Gamma_{\mathrm{top}} &:= \Omega \times \left\{\frac{t}{2}\right\}, \qquad \Gamma_{\mathrm{bot}} := \Omega \times \left\{-\frac{t}{2}\right\}, \\ \Gamma_{\mathrm{lat}} &:= \partial\Omega \times \left(-\frac{t}{2}, \frac{t}{2}\right). \end{aligned} \tag{6.6}$$

These sets cover $\partial\Omega^t$ up to edge curves of zero surface measure, which do not contribute to the boundary integrals.

Figure 6.1 summarizes the monolithic and panel-partitioned configurations used throughout the thermomechanical part of the thesis. The

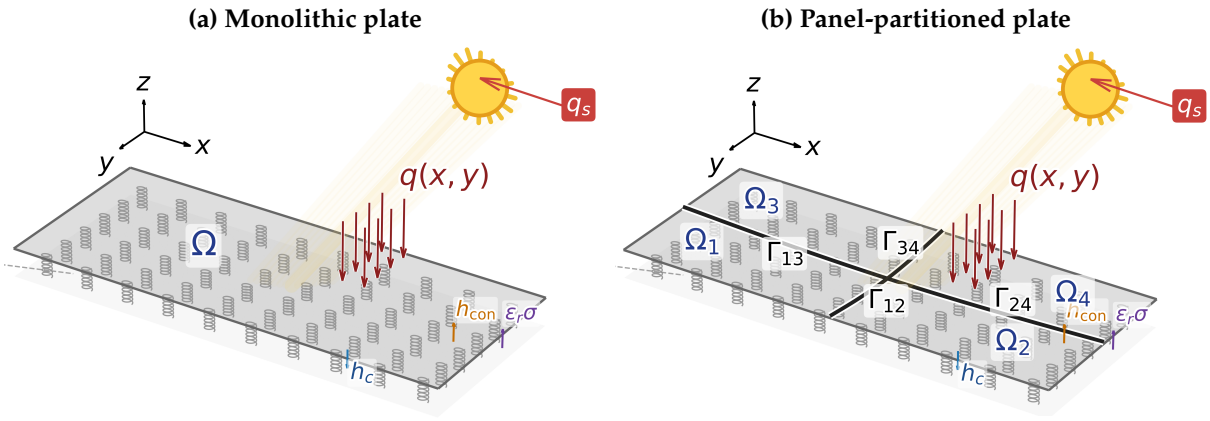



**Figure 6.1:** Geometric setting of the thermomechanical Kirchhoff–Love plate model. The plate is subjected to the transverse load $q(x, y)$ and absorbed solar flux $q_s$, with convection and long-wave radiation on the exposed face and displacement-dependent substrate exchange through $h_c(w)$. The panel-partitioned configuration retains distinct panel fields and couples adjacent panels across the physical interfaces $\Gamma_{ij}$.

diagram distinguishes the mechanical transverse load from the top-face environmental exchange and the displacement-dependent substrate exchange acting on the lower face.

## 6.3 Thermal subproblem

The temperature field is denoted by

$$T : \Omega^t \to \mathbb{R}.$$

Let $s : \Omega^t \to \mathbb{R}$ be a volumetric heat source, possibly $s \equiv 0$. The thermal conductivity is allowed to be diagonal and anisotropic:

$$\mathbf{K}_T := \begin{bmatrix} k_x & 0 & 0 \\ 0 & k_y & 0 \\ 0 & 0 & k_z \end{bmatrix}, \qquad k_x, k_y, k_z > 0. \tag{6.7}$$

The strong form in the body is

$$-\nabla \cdot (\mathbf{K}_T \nabla T) = s \qquad \text{in } \Omega^t. \tag{6.8}$$

With the outward-normal convention, the conductive heat flux leaving the body is

$$q_N := -\mathbf{N} \cdot \mathbf{K}_T \nabla T.$$

On the exposed top face, this outward heat flux balances convection to ambient air, long-wave radiation, and absorbed solar input [17, 18]:

$$-\mathbf{N} \cdot \mathbf{K}_T \nabla T = h_{\text{con}}(T - T_{\text{amb}}) + \varepsilon_{\text{r}} \sigma_{\text{SB}} \left(T^4 - T_{\text{amb}}^4\right) - q_s \qquad \text{on } \Gamma_{\text{top}}. \tag{6.9}$$

Here $h_{\text{con}}$ is the convection coefficient, $T_{\text{amb}}$ is the ambient temperature, $\varepsilon_{\text{r}} \in (0, 1]$ is the surface emissivity, $\sigma_{\text{SB}}$ is the Stefan–Boltzmann constant, and $q_s \geq 0$ is the absorbed solar heat flux. If $q_{\text{sol}}$ denotes the incident solar irradiance and $\varrho_{\text{sol}} \in [0, 1)$ the solar reflectivity, then

$$q_s = (1 - \varrho_{\text{sol}}) q_{\text{sol}}.$$

In the formulation below, $q_s$ is treated as a prescribed environmental

input. Solar-radiation datasets, such as PVGIS and ENEA solar-radiation tables, provide typical sources for specifying incident solar data [86, 87].

On the lower face, heat exchange with the supporting medium at temperature $T_{\mathrm{sub}}$ is modeled by a Robin condition,

$$-\mathbf{N}\cdot\mathbf{K}_T\nabla T = h_c(w)(T-T_{\mathrm{sub}}) \qquad \text{on } \Gamma_{\mathrm{bot}}. \tag{6.10}$$

The coefficient $h_c(w)$ is evaluated from the transverse displacement on the midsurface and then lifted to $\Gamma_{\mathrm{bot}}$ by the projection $(x,y,z)\mapsto(x,y)$. Values on panel-interface curves do not affect the surface integral.

The displacement convention is the one fixed in Chapter 3: $w>0$ denotes upward displacement. Hence $w\le 0$ corresponds to contact or compression into the support, while $w>0$ corresponds to uplift or detachment. The ideal contact-gap heat-exchange law is

$$h_c(w) = h_c^{\mathrm{cont}}\chi_{\{w\le 0\}} + h_c^{\mathrm{gap}}\chi_{\{w>0\}}, \qquad h_c^{\mathrm{cont}} \gg h_c^{\mathrm{gap}} \ge 0, \tag{6.11}$$

where $h_c^{\mathrm{cont}}$ is the thermal contact exchange coefficient and $h_c^{\mathrm{gap}}$ is the weaker exchange coefficient in the detached state [19, 20]. In the computations, this discontinuous law is replaced by the smooth regularization

$$h_c(w) = h_c^{\mathrm{cont}} + \big(h_c^{\mathrm{gap}} - h_c^{\mathrm{cont}}\big)S(w), \qquad S(w) := \frac{1}{2}\big(1+\tanh(\eta_c w)\big), \tag{6.12}$$

with $\eta_c>0$. Since $\eta_c w$ must be dimensionless, $\eta_c$ has units of inverse length when $w$ is dimensional. With the upward-positive convention, $S(w)\approx 0$ in contact and $S(w)\approx 1$ in uplift.

The lateral boundary is taken adiabatic:

$$-\mathbf{N}\cdot\mathbf{K}_T\nabla T = 0 \qquad \text{on } \Gamma_{\mathrm{lat}}. \tag{6.13}$$

For a prescribed displacement $w$, the heat problem is nonlinear in $T$ because of the radiative term on $\Gamma_{\mathrm{top}}$, while the coefficient $h_c(w)$ is fixed during that thermal solve.

## 6.4 Weak formulation of the thermal problem

Let

$$V_T := H^1(\Omega^t).$$

Multiplying (6.8) by a test function $\varphi\in V_T$, integrating by parts, and using (6.9)–(6.13), the weak thermal problem is: for a prescribed displacement $w$, find $T\in V_T$ such that, for all $\varphi\in V_T$,

$$\begin{aligned}
&\int_{\Omega^t}(\mathbf{K}_T\nabla T)\cdot\nabla\varphi\,\mathrm{d}V + \int_{\Gamma_{\mathrm{top}}} h_{\mathrm{con}}T\varphi\,\mathrm{d}\Gamma + \int_{\Gamma_{\mathrm{top}}}\varepsilon_{\mathrm{r}}\sigma_{\mathrm{SB}}T^4\varphi\,\mathrm{d}\Gamma \\
&\qquad + \int_{\Gamma_{\mathrm{bot}}} h_c(w)T\varphi\,\mathrm{d}\Gamma \\
&= \int_{\Omega^t} s\varphi\,\mathrm{d}V + \int_{\Gamma_{\mathrm{top}}} h_{\mathrm{con}}T_{\mathrm{amb}}\varphi\,\mathrm{d}\Gamma + \int_{\Gamma_{\mathrm{top}}}\varepsilon_{\mathrm{r}}\sigma_{\mathrm{SB}}T_{\mathrm{amb}}^4\varphi\,\mathrm{d}\Gamma \\
&\qquad + \int_{\Gamma_{\mathrm{top}}} q_s\varphi\,\mathrm{d}\Gamma + \int_{\Gamma_{\mathrm{bot}}} h_c(w)T_{\mathrm{sub}}\varphi\,\mathrm{d}\Gamma.
\end{aligned} \tag{6.14}$$

The top-face solar term appears on the right-hand side because $q_s$ is an incoming absorbed heat flux, while the convection and radiation terms are outward losses under the convention used in (6.9).

The thermal data are collected in the parameter vector

$$\mu_{\text{heat}} := \left(k_x, k_y, k_z, s, T_{\text{amb}}, T_{\text{sub}}, h_{\text{con}}, \varepsilon_{\text{r}}, q_s, h_c^{\text{cont}}, h_c^{\text{gap}}, \eta_c\right). \tag{6.15}$$

For a prescribed displacement $w$, the corresponding thermal solution operator is denoted by

$$T = \mathcal{T}_{\mu_{\text{heat}}}(w). \tag{6.16}$$

## 6.5 Thermal-to-plate reduction

The Kirchhoff–Love plate equation does not use the full three-dimensional temperature field directly. Thermal bending is driven by the part of the temperature distribution that is odd with respect to the centered thickness coordinate [13, 88]. Let $T_{\text{ref}}$ be a constant stress-free reference temperature and define

$$\Delta T := T - T_{\text{ref}}. \tag{6.17}$$

The plate thermal driver is the centered first through-thickness moment

$$\theta(x, y) := \frac{12}{t^3} \int_{-t/2}^{t/2} z\, \Delta T(x, y, z)\, \mathrm{d}z. \tag{6.18}$$

The prefactor is chosen so that, when $\Delta T(x, y, z) = z\, T^1(x, y)$, the reduction gives $\theta(x, y) = T^1(x, y)$. Since the thickness coordinate is centered,

$$\int_{-t/2}^{t/2} z\, \mathrm{d}z = 0,$$

and therefore $\theta$ is unaffected by the particular constant reference temperature used in $\Delta T$. The scalar field $\theta$ has the units of a through-thickness temperature gradient and is the only thermal field entering the plate bending equation. The thermal-to-plate reduction is written compactly as

$$\theta = \mathcal{M}(\Delta T). \tag{6.19}$$

A prescribed plate thermal driver can also be used directly. In that case, the heat problem (6.14) and the reduction (6.18) are bypassed, and $\theta$ is treated as known data in the mechanical problem below.

## 6.6 Thermal bending moments

The mechanical unknown is the panel-wise transverse displacement

$$w = \{w_i\}_{i=1}^{M}, \qquad w_i : \Omega_i \to \mathbb{R},$$

with the same upward-positive convention used in Chapter 3. No applied in-plane loading is considered in this thermomechanical formulation. On each panel, the Kirchhoff–Love curvature vector is

$$\boldsymbol{\kappa}(w_i) := \begin{bmatrix} \partial_{xx}^2 w_i \\ \partial_{yy}^2 w_i \\ 2\partial_{xy}^2 w_i \end{bmatrix}. \tag{6.20}$$

The orthotropic axes are assumed to be aligned with the Cartesian directions. The bending rigidity matrix is

$$\mathbf{D} \coloneqq \begin{bmatrix} D_x & D_{xy} & 0 \\ D_{xy} & D_y & 0 \\ 0 & 0 & D_s \end{bmatrix}, \tag{6.21}$$

where $D_x$ and $D_y$ are the principal bending rigidities, $D_{xy}$ is the bending-coupling rigidity, and $D_s$ is the twisting rigidity. The rigidities are assumed to define a positive bending energy on admissible curvature fields.

For aligned orthotropy, the plate thermal expansion vector is

$$\mathbf{B}^{\mathrm{th}} \coloneqq \begin{bmatrix} \alpha_x \\ \alpha_y \\ 0 \end{bmatrix}, \tag{6.22}$$

where $\alpha_x$ and $\alpha_y$ are the thermal expansion coefficients in the principal material directions. The thermal bending vector induced by the plate thermal driver is

$$\mathbf{m}^{\mathrm{th}}(\theta) \coloneqq \mathbf{D}\mathbf{B}^{\mathrm{th}}\theta. \tag{6.23}$$

Equivalently,

$$\mathbf{m}^{\mathrm{th}}(\theta) = \begin{bmatrix} (D_x\alpha_x + D_{xy}\alpha_y)\theta \\ (D_{xy}\alpha_x + D_y\alpha_y)\theta \\ 0 \end{bmatrix}. \tag{6.24}$$

Thus, in the aligned setting considered here, thermal bending contributes to the bending moments about the $x$- and $y$-directions, while no thermal twisting moment is introduced.

For a prescribed $\theta$, the bending contribution on $\Omega_i$ is

$$\begin{aligned} \mathcal{U}_{\mathrm{p},i}(w_i;\theta) = \frac{1}{2}\int_{\Omega_i} \boldsymbol{\kappa}(w_i)^{\top}\mathbf{D}\boldsymbol{\kappa}(w_i)\,\mathrm{d}\Omega \\ + \int_{\Omega_i} \mathbf{m}^{\mathrm{th}}(\theta)\cdot\boldsymbol{\kappa}(w_i)\,\mathrm{d}\Omega. \end{aligned} \tag{6.25}$$

This is the $w_i$-dependent part of the thermoelastic quadratic bending energy

$$\frac{1}{2}\int_{\Omega_i} \big(\boldsymbol{\kappa}(w_i) + \mathbf{B}^{\mathrm{th}}\theta\big)^{\top}\mathbf{D}\big(\boldsymbol{\kappa}(w_i) + \mathbf{B}^{\mathrm{th}}\theta\big)\,\mathrm{d}\Omega. \tag{6.26}$$

The remaining term,

$$\frac{1}{2}\int_{\Omega_i} (\mathbf{B}^{\mathrm{th}}\theta)^{\top}\mathbf{D}(\mathbf{B}^{\mathrm{th}}\theta)\,\mathrm{d}\Omega,$$

is independent of $w_i$ and therefore does not contribute to the mechanical stationarity condition.

## 6.7 Mechanical equilibrium with thermal loading

The unilateral Winkler-type foundation contribution on panel $\Omega_i$ is

$$\Psi_{\mathrm{f},i}(w_i) = \frac{1}{2}\int_{\Omega_i} k_s \left(\chi_{\{w_i<0\}} + \varepsilon_f \chi_{\{w_i\geq 0\}}\right) w_i^2 \,\mathrm{d}\Omega, \qquad 0 < \varepsilon_f \ll 1. \tag{6.27}$$

Here $k_s$ is the foundation stiffness and $\varepsilon_f$ is the uplift regularization factor. Under the upward-positive convention, the full stiffness acts in compression, while the uplift branch uses the reduced stiffness $\varepsilon_f k_s$. For prescribed $\theta$, this active compression-uplift branch is the only nonlinearity in the continuous mechanical subproblem.

Let $q$ be the total transverse load per unit area, positive upward. The external work on panel $\Omega_i$ is

$$\mathcal{W}_{\mathrm{ext},i}(w_i) = \int_{\Omega_i} q\, w_i \,\mathrm{d}\Omega. \tag{6.28}$$

When self-weight is included,

$$q = q_{\mathrm{mech}} + q_{\mathrm{self}}, \qquad q_{\mathrm{self}} = -\rho_m g t, \tag{6.29}$$

where $q_{\mathrm{mech}}$ is the prescribed mechanical load, $\rho_m$ is the panel density, and $g$ is the gravitational acceleration. If self-weight is neglected, then $q_{\mathrm{self}} = 0$ and $q = q_{\mathrm{mech}}$.

Mechanical transmission across panel interfaces is represented by the penalty energy

$$\Pi_{\mathrm{int}}(w) = \sum_{1\leq i<j\leq M} \frac{1}{2}\int_{\Gamma_{ij}} \left(\eta_w \left[\!\left[w\right]\!\right]_{ij}^2 + \eta_\psi \left[\!\left[\partial_{\mathbf{n}} w\right]\!\right]_{ij}^2\right) \mathrm{d}\Gamma. \tag{6.30}$$

The positive weights $\eta_w$ and $\eta_\psi$ control the weak enforcement of displacement continuity and normal-rotation continuity, respectively. The notation $\eta_\psi$ is used for the rotation penalty to avoid confusion with the thermal driver $\theta$. For a monolithic panel, $M = 1$, the interface sum is empty.

For prescribed $\theta$, the total potential energy is

$$\Pi(w;\theta) = \sum_{i=1}^{M} \left[\mathcal{U}_{\mathrm{p},i}(w_i;\theta) + \Psi_{\mathrm{f},i}(w_i) - \mathcal{W}_{\mathrm{ext},i}(w_i)\right] + \Pi_{\mathrm{int}}(w). \tag{6.31}$$

Mechanical equilibrium is obtained by stationarity of $\Pi(\cdot;\theta)$ with respect to admissible variations.

Let

$$V := \left\{ v = \{v_i\}_{i=1}^M \;\middle|\; v_i \in H^2(\Omega_i), \quad v_i = 0 \text{ on } \partial\Omega_i \cap \Gamma_S \right\}. \tag{6.32}$$

For a prescribed thermal driver $\theta$, the weak mechanical problem is: find $w \in V$ such that, for all $v \in V$,

$$a_{\mu_{\mathrm{mech}}}(w,v) + r^{\mathrm{th}}_{\mu_{\mathrm{mech}}}(\theta;v) + g_{\mu_{\mathrm{mech}}}(w;v) + c_{\mu_{\mathrm{mech}}}(w,v) = \ell_{\mu_{\mathrm{mech}}}(v). \tag{6.33}$$

The elastic bending form is

$$a_{\boldsymbol{\mu}_{\mathrm{mech}}}(w,v) = \sum_{i=1}^{M} \int_{\Omega_i} \Big[ D_x(\partial_{xx}^2 w_i)(\partial_{xx}^2 v_i) + D_y(\partial_{yy}^2 w_i)(\partial_{yy}^2 v_i) + D_{xy}\Big((\partial_{xx}^2 w_i)(\partial_{yy}^2 v_i) + (\partial_{yy}^2 w_i)(\partial_{xx}^2 v_i)\Big) + 4D_s(\partial_{xy}^2 w_i)(\partial_{xy}^2 v_i)\Big] \,\mathrm{d}\Omega. \tag{6.34}$$

The thermal bending functional is

$$r_{\boldsymbol{\mu}_{\mathrm{mech}}}^{\mathrm{th}}(\theta; v) = \sum_{i=1}^{M} \int_{\Omega_i} \big(\mathbf{D}\mathbf{B}^{\mathrm{th}}\theta\big) \cdot \boldsymbol{\kappa}(v_i) \,\mathrm{d}\Omega. \tag{6.35}$$

In components,

$$r_{\boldsymbol{\mu}_{\mathrm{mech}}}^{\mathrm{th}}(\theta; v) = \sum_{i=1}^{M} \int_{\Omega_i} \Big[ (D_x\alpha_x + D_{xy}\alpha_y)\theta\, \partial_{xx}^2 v_i + (D_{xy}\alpha_x + D_y\alpha_y)\theta\, \partial_{yy}^2 v_i \Big] \,\mathrm{d}\Omega. \tag{6.36}$$

For a given $\theta$, this functional is independent of $w$ and linear in the test function. It is kept on the left-hand side in (6.33) because it is the variation of the thermoelastic bending energy. Equivalently, it may be moved to the right-hand side as an external thermal bending load with the opposite sign.

The foundation contribution is

$$g_{\boldsymbol{\mu}_{\mathrm{mech}}}(w; v) = \sum_{i=1}^{M} \int_{\Omega_i} k_s \left(\chi_{\{w_i<0\}} + \varepsilon_f \chi_{\{w_i\geq 0\}}\right) w_i v_i \,\mathrm{d}\Omega. \tag{6.37}$$

The interface penalty form is

$$c_{\boldsymbol{\mu}_{\mathrm{mech}}}(w,v) = \sum_{1\leq i<j\leq M} \int_{\Gamma_{ij}} \Big( \eta_w \,[\![w]\!]_{ij}\,[\![v]\!]_{ij} + \eta_\psi \,[\![\partial_{\mathbf{n}} w]\!]_{ij}\,[\![\partial_{\mathbf{n}} v]\!]_{ij} \Big) \,\mathrm{d}\Gamma. \tag{6.38}$$

Finally, the load functional is

$$\ell_{\boldsymbol{\mu}_{\mathrm{mech}}}(v) = \sum_{i=1}^{M} \int_{\Omega_i} q\, v_i \,\mathrm{d}\Omega. \tag{6.39}$$

For smooth solutions, the physical total moment resultants associated with the elastic and thermal bending contributions are useful for interpreting boundary conditions. Suppressing the panel index, define

$$\begin{aligned} M_{xx}^{\mathrm{tot}} &= -\Big[D_x\partial_{xx}^2 w + D_{xy}\partial_{yy}^2 w + (D_x\alpha_x + D_{xy}\alpha_y)\theta\Big], \\ M_{yy}^{\mathrm{tot}} &= -\Big[D_{xy}\partial_{xx}^2 w + D_y\partial_{yy}^2 w + (D_{xy}\alpha_x + D_y\alpha_y)\theta\Big], \\ M_{xy}^{\mathrm{tot}} &= -2D_s\partial_{xy}^2 w. \end{aligned} \tag{6.40}$$

The total moment tensor is

$$\mathbf{M}^{\mathrm{tot}} := \begin{bmatrix} M_{xx}^{\mathrm{tot}} & M_{xy}^{\mathrm{tot}} \\ M_{xy}^{\mathrm{tot}} & M_{yy}^{\mathrm{tot}} \end{bmatrix}.$$

On a smooth exterior boundary segment,

$$M_n^{\mathrm{tot}} := \mathbf{n}^\top \mathbf{M}^{\mathrm{tot}} \mathbf{n}, \qquad M_{nt}^{\mathrm{tot}} := \mathbf{n}^\top \mathbf{M}^{\mathrm{tot}} \mathbf{t}. \tag{6.41}$$

The total Kirchhoff shear resultants are

$$\begin{aligned} Q_x^{\mathrm{tot}} &:= \partial_x M_{xx}^{\mathrm{tot}} + \partial_y M_{xy}^{\mathrm{tot}}, \\ Q_y^{\mathrm{tot}} &:= \partial_x M_{xy}^{\mathrm{tot}} + \partial_y M_{yy}^{\mathrm{tot}}, \end{aligned} \tag{6.42}$$

and

$$Q_n^{\mathrm{tot}} := Q_x^{\mathrm{tot}} n_x + Q_y^{\mathrm{tot}} n_y, \qquad V_n^{\mathrm{tot}} := Q_n^{\mathrm{tot}} + \partial_s M_{nt}^{\mathrm{tot}}. \tag{6.43}$$

The exterior boundary conditions are

$$\begin{cases} w = 0, \quad M_n^{\mathrm{tot}} = 0, & \text{on } \Gamma_S \quad \text{(simply supported)}, \\ M_n^{\mathrm{tot}} = 0, \quad V_n^{\mathrm{tot}} = 0, & \text{on } \Gamma_F \quad \text{(free edge)}. \end{cases} \tag{6.44}$$

The displacement condition on $\Gamma_S$ is essential; the moment and free-edge conditions are natural. Across an ideal panel interface, displacement and normal rotation are the kinematic transmission quantities, while the corresponding total normal moment and effective shear are balanced. In the present formulation, the kinematic transmission conditions are enforced weakly by (6.30); the associated equilibrium of interface resultants follows from the variational statement rather than from additional strong constraints.

The mechanical and thermoelastic data are collected in

$$\boldsymbol{\mu}_{\mathrm{mech}} := \left(D_x, D_y, D_{xy}, D_s, \alpha_x, \alpha_y, k_s, \varepsilon_f, q_{\mathrm{mech}}, \rho_m, \eta_w, \eta_\psi\right). \tag{6.45}$$

Together with the thermal parameters, the complete thermomechanical parameter vector is

$$\boldsymbol{\mu} := \left(\boldsymbol{\mu}_{\mathrm{mech}}, \boldsymbol{\mu}_{\mathrm{heat}}\right) \in \mathcal{D}, \tag{6.46}$$

where $\mathcal{D}$ is the admissible parameter set. For prescribed $\theta$, the mechanical solution operator is denoted by

$$w = \mathcal{K}_{\boldsymbol{\mu}_{\mathrm{mech}}}(\theta). \tag{6.47}$$

## 6.8 Partitioned staggered thermomechanical coupling

The thermal and mechanical subproblems are coupled through two mechanisms. First, the mechanical equation depends on the thermal solution through the thermal driver $\theta = \mathcal{M}(T - T_{\mathrm{ref}})$. Second, the thermal problem may depend on the current transverse displacement through the lower-face exchange coefficient $h_c(w)$. This produces a displacement–temperature feedback loop:

$$w \longmapsto h_c(w) \longmapsto T \longmapsto \theta \longmapsto w.$$

At the continuous level, a coupled solution is therefore a fixed point of the map

$$w = \mathcal{K}_{\boldsymbol{\mu}_{\mathrm{mech}}}\left(\mathcal{M}\left[\mathcal{T}_{\boldsymbol{\mu}_{\mathrm{heat}}}(w) - T_{\mathrm{ref}}\right]\right). \tag{6.48}$$

Given $\boldsymbol{\mu} \in \mathbb{D}$, a fixed reference temperature $T_{\mathrm{ref}}$, an initial displacement $w^{(0)}$, a relaxation parameter $0 < \omega \leq 1$, tolerances $\mathrm{tol}_w$ and $\mathrm{tol}_\theta$, a safeguard $\varepsilon_0 > 0$, and a maximum number of iterations $K_{\max}$, the partitioned staggered iteration proceeds as follows.

For $k = 0, 1, \ldots, K_{\max} - 1$, first solve the thermal problem with the current displacement:

$$T^{(k+1)} = \mathcal{T}_{\boldsymbol{\mu}_{\mathrm{heat}}}\left(w^{(k)}\right). \tag{6.49}$$

Then compute the temperature increment and the plate thermal driver:

$$\Delta T^{(k+1)} = T^{(k+1)} - T_{\mathrm{ref}}, \qquad \theta^{(k+1)} = \mathcal{M}\left(\Delta T^{(k+1)}\right). \tag{6.50}$$

With this thermal driver fixed, solve the mechanical subproblem:

$$\widetilde{w}^{(k+1)} = \mathcal{K}_{\boldsymbol{\mu}_{\mathrm{mech}}}\left(\theta^{(k+1)}\right). \tag{6.51}$$

The displacement is then relaxed according to

$$w^{(k+1)} = (1 - \omega) w^{(k)} + \omega \widetilde{w}^{(k+1)}. \tag{6.52}$$

In the fully discrete solver of Chapter 7, convergence of the iteration is monitored through relative finite element $L^2(\Omega)$ updates of the displacement and the plate thermal driver. The stopping criterion is evaluated after the relaxation update. For $k \geq 1$, the iteration is stopped if

$$\begin{aligned} \frac{\|w^{(k+1)} - w^{(k)}\|_{L^2(\Omega)}}{\max\{\|w^{(k+1)}\|_{L^2(\Omega)}, \varepsilon_0\}} &< \mathrm{tol}_w, \\ \frac{\|\theta^{(k+1)} - \theta^{(k)}\|_{L^2(\Omega)}}{\max\{\|\theta^{(k+1)}\|_{L^2(\Omega)}, \varepsilon_0\}} &< \mathrm{tol}_\theta. \end{aligned} \tag{6.53}$$

Here $\varepsilon_0 > 0$ prevents division by a vanishing field norm. If the stopping criterion is not satisfied within $K_{\max}$ iterations, the final iterate is marked as nonconverged. Such a state may be retained for solver diagnostics but is excluded from any database representing the converged coupled solution map.

**Thermomechanical feedback**

The partitioned formulation separates the nonlinearities by subproblem. Radiation is handled inside the thermal solve, the unilateral foundation is handled inside the mechanical solve, and displacement-dependent heat exchange is handled by the fixed-point iteration. The partitioned iteration is therefore not a monolithic thermomechanical Newton method.

## One-way specialization

A non-iterative one-way approximation is obtained from the preceding algorithm by prescribing

$$K_{\max} = 1, \qquad \omega = 1.$$

Starting from the prescribed initial or frozen field $w^{(0)}$, the thermal problem is solved once, the corresponding thermal driver is reconstructed,

and one mechanical equilibrium problem is evaluated:

$$\begin{aligned} T^{(1)} &= \mathcal{T}_{\mu_{\text{heat}}}\left(w^{(0)}\right), \\ \theta^{(1)} &= \mathcal{M}\left(T^{(1)} - T_{\text{ref}}\right), \\ w^{(1)} &= \mathcal{K}_{\mu_{\text{mech}}}\left(\theta^{(1)}\right). \end{aligned} \tag{6.54}$$

The returned displacement is not supplied to a second thermal solve. This specialization suppresses feedback from the computed mechanical state to the lower-face exchange coefficient and is appropriate only when that feedback is negligible for the parameter regime under consideration.

## 6.9 Parameter dependence and feedback structure

The parameter dependence of the thermomechanical model separates naturally into mechanical and thermal blocks. The mechanical block $\mu_{\text{mech}}$ contains bending rigidities, thermal expansion coefficients, foundation and interface parameters, mechanical loading, and density. The thermal block $\mu_{\text{heat}}$ contains conductivity, source and environmental data, radiative and convective parameters, and the contact-gap heat-exchange parameters. The reference temperature $T_{\text{ref}}$ is used to form the temperature increment, but it does not affect the continuous value of $\theta$ when it is constant.

The accepted partitioned state defines the parameterized thermomechanical solution map

$$\mu \longmapsto \big(T(\mu), \theta(\mu), w(\mu)\big). \tag{6.55}$$

The one-way approximation in Eq. (6.54) is the single-iteration specialization of the partitioned procedure. The two-way model instead retains the fixed-point dependence defined in Section 6.8 and Eq. (6.48). Consequently, only states that satisfy the coupled stopping criterion represent accepted evaluations of the solution map in Eq. (6.55); an unconverged iterate is not used as a snapshot of that map.

Chapter 7 replaces the continuous operators $\mathcal{T}_{\mu_{\text{heat}}}$, $\mathcal{M}$, and $\mathcal{K}_{\mu_{\text{mech}}}$ by their finite element counterparts. The accepted FOM states then provide the two outputs reduced separately in Chapter 8: global displacement and global thermal driver.

## 6.10 Summary of the thermomechanical formulation

The thermomechanical formulation extends the mechanical plate framework developed in Part I. The thermal subproblem is a steady three-dimensional heat-conduction problem on $\Omega^t$, with top-face convection, long-wave radiation, absorbed solar loading, displacement-dependent lower-face exchange, and adiabatic lateral boundaries. The lower-face exchange may depend on transverse displacement through a smooth contact–gap law, introducing thermomechanical feedback.

The three-dimensional temperature field enters the plate equation only through the centered first through-thickness moment $\theta$. This reduced

thermal driver produces thermal bending moments $\mathbf{m}^{\mathrm{th}}(\theta) = \mathbf{D}\mathbf{B}^{\mathrm{th}}\theta$, which modify the Kirchhoff–Love mechanical equilibrium while preserving the mechanical sign convention, foundation model, and panel-interface structure used earlier in the thesis.

The resulting thermomechanical model is resolved by a partitioned two-way iteration that alternates the thermal problem, thermal-to-plate reduction, and mechanical equilibrium until both displacement and thermal-driver updates satisfy the prescribed stopping criterion. The one-way approximation is recovered by one unrelaxed pass through the same sequence. Chapter 7 constructs the corresponding high-fidelity finite element approximation and examines its spatial refinement, iteration convergence, and computational cost.

# Full-Order Thermomechanical Approximation and Coupled-Solver Assessment

# 7


7.1 Role and scope of the thermomechanical high-fidelity model . . . . . 74
7.2 Finite element approximation of the thermal subproblem . . . . . 74
7.3 Discrete reconstruction of the plate thermal driver . . 75
7.4 $C^0$ interior-penalty plate approximation with thermal bending . . . . . 76
7.5 Partitioned thermomechanical full-order solves . . . . . 78
7.6 Navier-series reference for the simply supported auxiliary problem . . . . . 79
7.7 Verification and convergence studies . . . . . 79
7.8 Summary of the thermomechanical full-order model . 86


## 7.1 Role and scope of the thermomechanical high-fidelity model

Chapter 6 introduced a continuous thermomechanical model in which a steady three-dimensional heat problem supplies a plate thermal driver and the plate displacement may, in turn, modify the lower-face heat exchange. The present chapter constructs the corresponding high-fidelity finite element approximation. The discretization must resolve four distinct operations without conflating them: the three-dimensional thermal solve, the through-thickness reduction, the fourth-order Kirchhoff–Love plate solve, and the thermomechanical fixed-point iteration.

Two independent meshes are used. A conforming tetrahedral mesh resolves the thermal domain $\Omega^t$, while a panel-aligned triangular mesh resolves the plate midsurface $\Omega$. The plate approximation inherits the symmetric $C^0$ interior-penalty construction used in Chapter 4, but it is augmented by the thermal bending functional of Section 6.7. Numerical stabilization on element facets internal to a panel remains distinct from the physical penalty coupling across panel interfaces.

The verification program comprises an auxiliary Navier-series comparison for a single-panel constant-coefficient problem, a plate-only refinement study that isolates the $C^0$ interior-penalty discretization, and a coupled refinement study that also examines convergence of the partitioned iteration. The Navier field and the refined finite element fields serve different purposes and are identified accordingly: the former is a finitely evaluated analytical series for a simplified constant-coefficient problem, whereas the latter are numerical reference solutions for the stated refinement tests.

## 7.2 Finite element approximation of the thermal subproblem

Let $\mathcal{T}_h^T$ be a shape-regular tetrahedral mesh of the three-dimensional plate domain $\Omega^t$. For polynomial degree $p_T \geq 1$, define the continuous Lagrange space

$$V_{T,h} := \left\{ \varphi_h \in C^0(\overline{\Omega^t}) \;\middle|\; \varphi_h|_K \in \mathbb{P}_{p_T}(K) \quad \forall K \in \mathcal{T}_h^T \right\}. \tag{7.1}$$

The boundary facets are marked according to the decomposition $\partial\Omega^t = \Gamma_{\text{top}} \cup \Gamma_{\text{bot}} \cup \Gamma_{\text{lat}}$, up to the measure-zero edge curves identified in Section 6.2.

The displacement supplied to the thermal subproblem is defined on the midsurface. Its value in the lower-face exchange coefficient is evaluated through the projection

$$\pi : \Gamma_{\text{bot}} \to \Omega, \qquad \pi(x, y, -t/2) = (x, y). \tag{7.2}$$

Thus, the Robin coefficient on $\Gamma_{\mathrm{bot}}$ is $h_c(w_h \circ \pi)$. For a panel-wise plate field, the traces on panel-interface curves do not alter the lower-face surface integrals because those curves have zero surface measure.

For a prescribed displacement $w_h$, the discrete thermal problem is: find $T_h \in V_{T,h}$ such that, for every $\varphi_h \in V_{T,h}$,

$$\begin{aligned} 0 = &\int_{\Omega^t} (\mathbf{K}_T \nabla T_h) \cdot \nabla \varphi_h \, \mathrm{d}V + \int_{\Gamma_{\mathrm{top}}} h_{\mathrm{con}} T_h \varphi_h \, \mathrm{d}\Gamma + \int_{\Gamma_{\mathrm{top}}} \varepsilon_{\mathrm{r}} \sigma_{\mathrm{SB}} T_h^4 \varphi_h \, \mathrm{d}\Gamma \\ &+ \int_{\Gamma_{\mathrm{bot}}} h_c(w_h \circ \pi) T_h \varphi_h \, \mathrm{d}\Gamma - \int_{\Omega^t} s \varphi_h \, \mathrm{d}V - \int_{\Gamma_{\mathrm{top}}} h_{\mathrm{con}} T_{\mathrm{amb}} \varphi_h \, \mathrm{d}\Gamma \\ &- \int_{\Gamma_{\mathrm{top}}} \varepsilon_{\mathrm{r}} \sigma_{\mathrm{SB}} T_{\mathrm{amb}}^4 \varphi_h \, \mathrm{d}\Gamma - \int_{\Gamma_{\mathrm{top}}} q_s \varphi_h \, \mathrm{d}\Gamma - \int_{\Gamma_{\mathrm{bot}}} h_c(w_h \circ \pi) T_{\mathrm{sub}} \varphi_h \, \mathrm{d}\Gamma. \end{aligned} \tag{7.3}$$

During one thermal solve, $h_c(w_h \circ \pi)$ is evaluated from the supplied displacement and held fixed. The thermal algebraic problem remains nonlinear when $\varepsilon_{\mathrm{r}} > 0$ because of the radiative term. The implementation used in this study treats this thermal nonlinearity by Newton's method [89].

## 7.3 Discrete reconstruction of the plate thermal driver

Let $\mathcal{T}_h^{\Omega}$ be the triangular plate mesh introduced below. For degree $p_\theta \geq 1$, define

$$\Theta_h := \left\{ \psi_h \in C^0(\overline{\Omega}) \;\middle|\; \psi_h|_K \in \mathbb{P}_{p_\theta}(K) \quad \forall K \in \mathcal{T}_h^{\Omega} \right\}. \tag{7.4}$$

After solving Eq. (7.3), the reference temperature is represented in $V_{T,h}$, the increment $\Delta T_h = T_h - T_{\mathrm{ref}}$ is formed, and the centered first through-thickness moment is interpolated into $\Theta_h$:

$$\theta_h = I_{\Theta,h}\left[ \frac{12}{t^3} \int_{-t/2}^{t/2} z \, \Delta T_h(\cdot, z) \, \mathrm{d}z \right]. \tag{7.5}$$

Here $I_{\Theta,h}$ denotes nodal interpolation into $\Theta_h$. The thickness integral is evaluated numerically, and $\theta_h$ has units of temperature per unit length, $\mathrm{K\,m^{-1}}$. The reduction formula is fixed, whereas the number of through-thickness quadrature points used to evaluate its integral is a study-specific numerical setting. The value used in each verification and ROM campaign is therefore reported with that study.

The resulting discrete thermal and reduction maps are denoted by

$$T_h = \mathcal{T}_{h,\boldsymbol{\mu}_{\mathrm{heat}}}(w_h), \qquad \theta_h = \mathcal{M}_h(\Delta T_h). \tag{7.6}$$

The two mappings act between different finite element spaces and need not use matching meshes or polynomial degrees [90].

## 7.4 $C^0$ interior-penalty plate approximation with thermal bending

### Panel-wise finite element space

The plate triangulation $\mathcal{T}_h^\Omega$ is shape regular and aligned with the panel partition: each physical interface $\Gamma_{ij}$ is a union of mesh facets. Its restriction to panel $\Omega_i$ is $\mathcal{T}_h^i$. For degree $p \geq 2$, define

$$\begin{aligned} W_h^i &:= \left\{ \phi_h \in C^0(\overline{\Omega}_i) \;\middle|\; \phi_h|_K \in \mathbb{P}_p(K) \quad \forall K \in \mathcal{T}_h^i \right\}, \\ W_h &:= \bigoplus_{i=1}^{M} W_h^i. \end{aligned} \tag{7.7}$$

The direct sum keeps the degrees of freedom belonging to adjacent physical panels distinct before the interface terms are applied. In the implementation inherited from the mechanical full-order formulation, the mixed finite element function contains one scalar component $w_{h,i}$ for each panel. All components are stored on the same panel-aligned mesh, but $w_{h,i}$ represents the physical displacement only on $\Omega_i$. The physical piecewise field is therefore recovered from the corresponding component on each panel,

$$w_h|_{\Omega_i} = w_{h,i}|_{\Omega_i}, \qquad w_h = \sum_{i=1}^{M} \chi_{\Omega_i} w_{h,i}.$$

The admissible discrete space is

$$V_h := \left\{ \phi_h \in W_h \;\middle|\; \phi_h = 0 \text{ on } \Gamma_S \right\}. \tag{7.8}$$

When the exterior edge is free, $\Gamma_S = \emptyset$ and no essential displacement condition is imposed.

### Numerical interior-facet stabilization

The continuous elastic, thermal, foundation, interface, and load forms of Chapter 6 are evaluated elementwise. The additional term required by the $C^0$ approximation acts only on numerical facets lying strictly inside each physical panel.

Let $E = \partial K^+ \cap \partial K^-$ be an interior facet with neighboring elements $K^+, K^- \in \mathcal{T}_h^i$. Fix a unit normal $\boldsymbol{\nu}_E$ directed from $K^+$ to $K^-$, let $h_E$ be the average diameter of the two elements, and denote traces by superscripts $+$ and $-$. For a scalar or tensor quantity $\zeta$, define

$$\{\{\zeta\}\}_E := \frac{1}{2}(\zeta^+ + \zeta^-), \qquad \left[\left[\partial_{\boldsymbol{\nu}_E} \phi_h\right]\right]_E := \partial_{\boldsymbol{\nu}_E} \phi_h^+ - \partial_{\boldsymbol{\nu}_E} \phi_h^-. \tag{7.9}$$

These operators belong to the numerical interior-penalty formulation and are distinct from the physical panel-interface jumps of Eq. (6.4).

Define the effective bending scale and facet penalty by

$$D_{\text{eff}} := \sqrt{D_x D_y}, \qquad \tau_E := \frac{(p+2)^2 D_{\text{eff}}}{h_E}. \tag{7.10}$$

For $\phi_h \in V_h$, introduce the elastic bending tensor

$$\mathbf{B}^{\mathrm{el}}(\phi_h) \coloneqq \begin{bmatrix} D_x \partial^2_{xx}\phi_h + D_{xy}\partial^2_{yy}\phi_h & 2D_s\partial^2_{xy}\phi_h \\ 2D_s\partial^2_{xy}\phi_h & D_{xy}\partial^2_{xx}\phi_h + D_y\partial^2_{yy}\phi_h \end{bmatrix}, \tag{7.11}$$

with normal elastic moment

$$m_E^{\mathrm{el}}(\phi_h) \coloneqq \boldsymbol{\nu}_E^{\top}\mathbf{B}^{\mathrm{el}}(\phi_h)\boldsymbol{\nu}_E. \tag{7.12}$$

The tensor in Eq. (7.11) is associated with the elastic bilinear form. It differs by the conventional leading sign from the physical total moments in Eq. (6.40).

Let $\mathscr{E}_{h,i}^{\mathrm{int}}$ denote the numerical facets strictly internal to panel $\Omega_i$. The symmetric $C^0$ interior-penalty contribution is

$$\begin{aligned} s_h(w_h, \phi_h) \coloneqq \sum_{i=1}^{M} \sum_{E \in \mathscr{E}_{h,i}^{\mathrm{int}}} \Bigg[ & - \int_E \{\{m_E^{\mathrm{el}}(w_h)\}\}_E \left[\left[\partial_{\boldsymbol{\nu}_E}\phi_h\right]\right]_E \, \mathrm{d}s \\ & - \int_E \{\{m_E^{\mathrm{el}}(\phi_h)\}\}_E \left[\left[\partial_{\boldsymbol{\nu}_E} w_h\right]\right]_E \, \mathrm{d}s \\ & + \int_E \tau_E \left[\left[\partial_{\boldsymbol{\nu}_E} w_h\right]\right]_E \left[\left[\partial_{\boldsymbol{\nu}_E}\phi_h\right]\right]_E \, \mathrm{d}s \Bigg]. \end{aligned} \tag{7.13}$$

This term weakly restores the interelement normal-slope continuity required by the fourth-order plate operator. It is a numerical stabilization and not a constitutive model for a physical panel seam [25–27, 38].

## Physical panel-interface coupling and discrete residual

For a physical interface facet $E \subset \Gamma_{ij}$, let $h_E$ be the average diameter of the adjacent plate elements. The reported computations use the facetwise interface weights

$$\eta_{w,E} = p\frac{D_{\mathrm{eff}}}{h_E^3}, \qquad \eta_{\psi,E} = \frac{12^{p+1}}{p}\frac{D_{\mathrm{eff}}}{h_E}. \tag{7.14}$$

The corresponding discrete physical interface form is

$$\begin{aligned} c_{\mu_{\mathrm{mech}},h}(w_h, \phi_h) \coloneqq \sum_{1 \le i < j \le M} \int_{\Gamma_{ij}} \Big( & \eta_{w,E} \left[\!\left[ w_h \right]\!\right]_{ij} \left[\left[\phi_h\right]\right]_{ij} \\ & + \eta_{\psi,E} \left[\!\left[\partial_{\mathbf{n}} w_h\right]\!\right]_{ij} \left[\left[\partial_{\mathbf{n}}\phi_h\right]\right]_{ij} \Big) \, \mathrm{d}\Gamma. \end{aligned} \tag{7.15}$$

The weights in Eq. (7.14) are distinct from $\tau_E$. They are fixed by the discretization in the numerical studies below and are not independently sampled coordinates in the thermomechanical ROM databases.

For prescribed $\theta_h \in \Theta_h$, the discrete mechanical problem is: find $w_h \in V_h$ such that

$$\mathscr{F}_h(w_h; \theta_h, \phi_h) = 0 \qquad \forall \phi_h \in V_h. \tag{7.16}$$

The elastic bending, thermal bending, foundation, and load forms are defined in Eqs. (6.34), (6.35), (6.37), and (6.39), respectively. The numerical interior-facet form $s_h$ is defined in Eq. (7.13), and the discrete physical panel-interface form $c_{\mu_{\mathrm{mech}},h}$ is defined in Eq. (7.15). With these

definitions,

$$\begin{aligned}\mathcal{F}_h(w_h;\theta_h,\phi_h) \coloneqq\ & a_{\boldsymbol{\mu}_{\mathrm{mech}}}(w_h,\phi_h) + s_h(w_h,\phi_h) + r^{\mathrm{th}}_{\boldsymbol{\mu}_{\mathrm{mech}}}(\theta_h;\phi_h) \\ & + g_{\boldsymbol{\mu}_{\mathrm{mech}}}(w_h;\phi_h) + c_{\boldsymbol{\mu}_{\mathrm{mech}},h}(w_h,\phi_h) - \ell_{\boldsymbol{\mu}_{\mathrm{mech}}}(\phi_h).\end{aligned} \tag{7.17}$$

For fixed $\theta_h$, the thermal bending term is independent of $w_h$. The unilateral foundation retains the active compression–uplift response of Eq. (6.37); the resulting nonlinear algebraic problem is solved by Newton's method. The discrete mechanical solution map is denoted by

$$w_h = \mathcal{K}_{h,\boldsymbol{\mu}_{\mathrm{mech}}}(\theta_h). \tag{7.18}$$

## 7.5 Partitioned thermomechanical full-order solves

The full-order solver replaces the continuous maps of Chapter 6 by

$$\mathcal{T}_{\boldsymbol{\mu}_{\mathrm{heat}}} \mapsto \mathcal{T}_{h,\boldsymbol{\mu}_{\mathrm{heat}}}, \qquad \mathcal{M} \mapsto \mathcal{M}_h, \qquad \mathcal{K}_{\boldsymbol{\mu}_{\mathrm{mech}}} \mapsto \mathcal{K}_{h,\boldsymbol{\mu}_{\mathrm{mech}}}. \tag{7.19}$$

Starting from $w_h^{(0)}$, iteration $k$ computes

$$\begin{aligned} T_h^{(k+1)} &= \mathcal{T}_{h,\boldsymbol{\mu}_{\mathrm{heat}}}(w_h^{(k)}), \\ \theta_h^{(k+1)} &= \mathcal{M}_h(T_h^{(k+1)} - T_{\mathrm{ref}}), \\ \widetilde{w}_h^{(k+1)} &= \mathcal{K}_{h,\boldsymbol{\mu}_{\mathrm{mech}}}(\theta_h^{(k+1)}), \\ w_h^{(k+1)} &= (1-\omega)w_h^{(k)} + \omega\widetilde{w}_h^{(k+1)}. \end{aligned} \tag{7.20}$$

Only the displacement is relaxed; $\theta_h^{(k+1)}$ is the thermal driver reconstructed in the current thermal solve. The stopping test is the finite element $L^2(\Omega)$ criterion of Eq. (6.53), evaluated after displacement relaxation. If two fields are represented in different finite element spaces, the previous field is transferred to the current comparison space before forming the update.

If the test is satisfied at iteration $k$, the returned state is

$$\left(w_h^{(k+1)}, T_h^{(k+1)}, \theta_h^{(k+1)}\right). \tag{7.21}$$

The thermal fields in Eq. (7.21) were generated from $w_h^{(k)}$, whereas the accepted displacement is the relaxed update $w_h^{(k+1)}$. The remaining mismatch is controlled by the update tolerances. Failure to satisfy the stopping criterion within $K_{\max}$ iterations marks the final iterate as nonconverged.

The one-way full-order approximation is the specialization $K_{\max} = 1$ and $\omega = 1$ of Eq. (7.20). It performs one thermal solve from the prescribed initial displacement, one thermal-to-plate reduction, and one mechanical solve. The returned displacement is not supplied to a second evaluation of $h_c$.

## 7.6 Navier-series reference for the simply supported auxiliary problem

The first verification uses the classical Navier construction for a simply supported rectangular, single-panel, aligned-orthotropy plate. The construction applies to a prescribed spatially constant Winkler coefficient and a fixed plate thermal driver. It does not solve the panel-partitioned problem, the pointwise unilateral foundation law, or the displacement-dependent coupled thermal problem.

Let $\Omega = (0, a) \times (0, b)$ and define

$$\widehat{D}_{12} \coloneqq D_{xy} + 2D_s, \qquad \lambda_m = \frac{m\pi}{a}, \qquad \mu_n = \frac{n\pi}{b}, \qquad m, n \in \mathbb{N}. \tag{7.22}$$

For a prescribed constant effective foundation coefficient $k_s^{\mathrm{eff}}$, the auxiliary strong equation is

$$\begin{aligned} &D_x \partial^4_{xxxx} w + 2\widehat{D}_{12} \partial^4_{xxyy} w + D_y \partial^4_{yyyy} w + k_s^{\mathrm{eff}} w \\ &\quad = q - \partial^2_{xx} M^{\mathrm{th}}_{xx} - \partial^2_{yy} M^{\mathrm{th}}_{yy} \qquad \text{in } \Omega. \end{aligned} \tag{7.23}$$

with $w = 0$ and vanishing total normal bending moment on each boundary edge. Expanding $w$, $q$, and $\theta$ in the double-sine basis gives

$$w(x, y) = \sum_{m=1}^{\infty} \sum_{n=1}^{\infty} W_{mn} \sin(\lambda_m x) \sin(\mu_n y), \tag{7.24}$$

where

$$W_{mn} = \frac{q_{mn} + \left[ (D_x \alpha_x + D_{xy} \alpha_y) \lambda_m^2 + (D_{xy} \alpha_x + D_y \alpha_y) \mu_n^2 \right] \theta_{mn}}{D_x \lambda_m^4 + 2\widehat{D}_{12} \lambda_m^2 \mu_n^2 + D_y \mu_n^4 + k_s^{\mathrm{eff}}}. \tag{7.25}$$

The sine coefficients are

$$\begin{aligned} q_{mn} &= \frac{4}{ab} \int_0^a \int_0^b q(x, y) \sin(\lambda_m x) \sin(\mu_n y) \, \mathrm{d}y \, \mathrm{d}x, \\ \theta_{mn} &= \frac{4}{ab} \int_0^a \int_0^b \theta(x, y) \sin(\lambda_m x) \sin(\mu_n y) \, \mathrm{d}y \, \mathrm{d}x. \end{aligned} \tag{7.26}$$

The positive sign of the thermal contribution in Eq. (7.25) follows from the total-moment convention in Eq. (6.40). The finite series used below is denoted by $w_{\mathrm{ext}}$ to emphasize that it is a numerically evaluated reference field, not an exact solution of the complete model [13].

## 7.7 Verification and convergence studies

The computations were carried out with FEniCS/DOLFIN [37] for finite element assembly and data structures, UFL for the variational forms, and PETSc through the DOLFIN linear-algebra interface. The nonlinear algebraic problems associated with radiation and the unilateral foundation were treated by Newton's method. The wall-clock measurements reported below were obtained on a MacBook equipped with an Apple M5 processor with 10 cores (4 performance and 6 efficiency cores) and 16 GB of memory, running macOS 26.5.2. The recorded software versions were Python 3.12.13, DOLFIN 2019.1.0, UFL 2019.1.0, PETSc/petsc4py 3.21.2,

and mpi4py 3.1.6. The FEniCS process used the `x86_64` runtime through Rosetta 2. Each profiling run used one MPI rank, and the standard OpenMP and BLAS thread-count variables were set to one. Wall time was measured with the monotonic Python `time.perf_counter` clock. The reported timings are implementation- and platform-specific measurements rather than hardware-independent complexity estimates.

The common plate geometry is

$$\Omega = (0, L) \times (0, B), \qquad L = 2\ \mathrm{m}, \qquad B = 1\ \mathrm{m}, \qquad t = 0.05\ \mathrm{m}. \tag{7.27}$$

The verification cases use a simply supported exterior boundary, $\Gamma_S = \partial\Omega$ and $\Gamma_F = \emptyset$.

## Auxiliary Navier-series comparison

The mechanical data are

$$\begin{aligned} &D_x = D_y = 10^4\ \mathrm{N\,m}, && D_{xy} = 0.3 \times 10^4\ \mathrm{N\,m}, \\ &D_s = 0.35 \times 10^4\ \mathrm{N\,m}, && k_s = 10^3\ \mathrm{N\,m^{-3}}, \\ &q_{\mathrm{mech}} = -10^4\ \mathrm{N\,m^{-2}}, && \varepsilon_f = 10^{-3}. \end{aligned}$$

The Navier construction uses the prescribed constant value

$$k_s^{\mathrm{eff}} = k_s = 10^3\ \mathrm{N\,m^{-3}}. \tag{7.28}$$

The reconstructed displacement is nonnegative on $\Omega$. Therefore, Eq. (7.28) is not interpreted as a self-consistent compression branch of the unilateral law; it defines only the auxiliary constant-coefficient problem used for this comparison.

The thermal and thermoelastic data are

$$\begin{aligned} &T_{\mathrm{amb}} = 300.15\ \mathrm{K}, && T_{\mathrm{sub}} = 297.15\ \mathrm{K}, \\ &h_{\mathrm{con}} = 10\ \mathrm{W\,m^{-2}K^{-1}}, && \varepsilon_{\mathrm{r}} = 0.90, \\ &q_s = 700\ \mathrm{W\,m^{-2}}, && h_c^{\mathrm{cont}} = 200\ \mathrm{W\,m^{-2}K^{-1}}, \\ &h_c^{\mathrm{gap}} = 5\ \mathrm{W\,m^{-2}K^{-1}}, && \eta_c = 50\ \mathrm{m^{-1}}, \\ &k_x = k_y = k_z = 0.35\ \mathrm{W\,m^{-1}K^{-1}}, \\ &\alpha_x = \alpha_y = 1.3 \times 10^{-4}\ \mathrm{K^{-1}}, && \rho_m = 1050\ \mathrm{kg\,m^{-3}}. \end{aligned}$$

The plate mesh uses resolution parameter $N_\Omega = 64$ and degree $p = 2$. The heat mesh is generated from a $64 \times 32 \times 16$ box partition and tetrahedral subdivision with $p_T = 1$. The thermal-driver space uses $p_\theta = 3$, and Eq. (7.5) is evaluated with 20 through-thickness quadrature points.

The partitioned computation uses

$$\omega = 0.7, \qquad \mathrm{tol}_w = \mathrm{tol}_\theta = 10^{-6}, \qquad K_{\max} = 40.$$

It satisfies the stopping criterion after 12 iterations. The final relative updates are $8.25{\times}10^{-7}$ for $w_h$ and $9.88{\times}10^{-9}$ for $\theta_h$. The constant reference temperature is the volume average of the computed temperature field,

$$T_{\mathrm{ref}} = 314.600821\ \mathrm{K}.$$

As established in Section 6.5, this constant does not affect the continuous centered first moment.

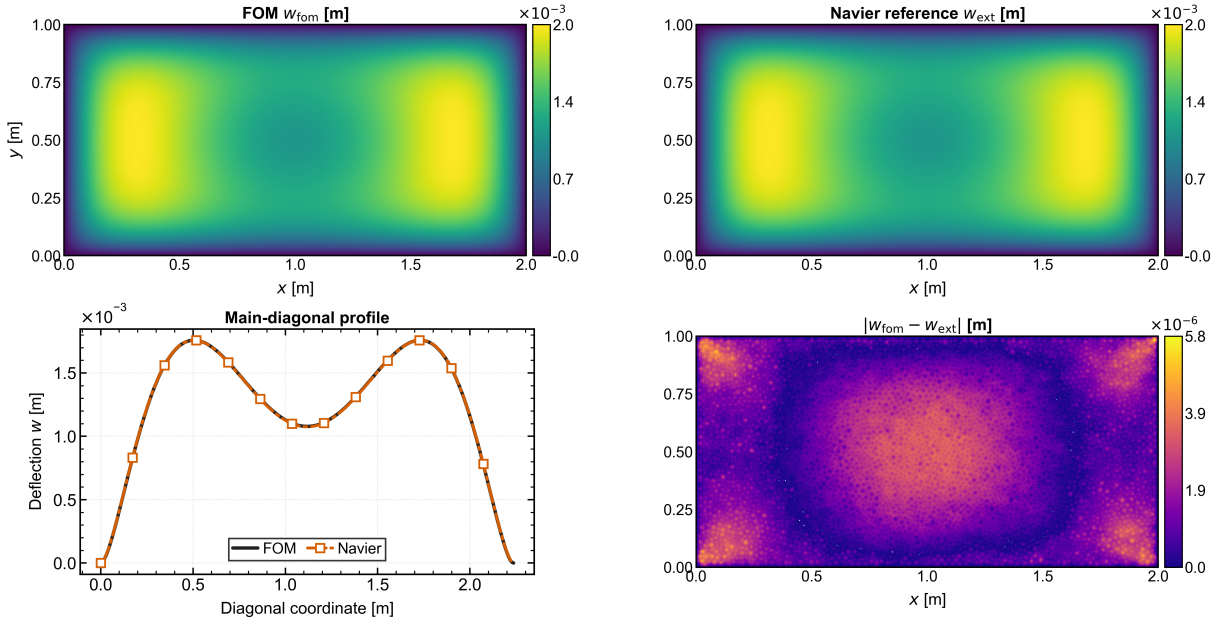



**Figure 7.1:** Auxiliary Navier-series comparison for the simply supported, single-panel constant-coefficient problem. The top row shows the finite element displacement $w_h$ (left) and the finitely evaluated Navier-series reference $w_{\text{ext}}$ (right). The bottom-left panel compares the two fields along the plate diagonal, with open square markers identifying the Navier reference, while the bottom-right panel shows the pointwise absolute difference $|w_h - w_{\text{ext}}|$.

The spatially varying thermal driver is projected onto the sine basis using a $401 \times 401$ in-plane grid. The finite evaluation retains $1 \le m, n \le 159$, corresponding to the implementation setting of 80 nominal series terms, and the reconstructed reference field is represented with polynomial degree 8.

Figure 7.1 compares the finite element and Navier-series solutions through full-field contours, a main-diagonal profile, and the pointwise absolute difference. The two solutions reproduce the same symmetric two-lobe displacement pattern, and their main-diagonal profiles are indistinguishable at the plotted scale. The displacement is of order $10^{-3}$ m, whereas the pointwise difference remains below approximately $5.8 \times 10^{-6}$ m. This agreement supports the finite element assembly of the elastic and thermal-bending contributions for the stated single-panel problem.

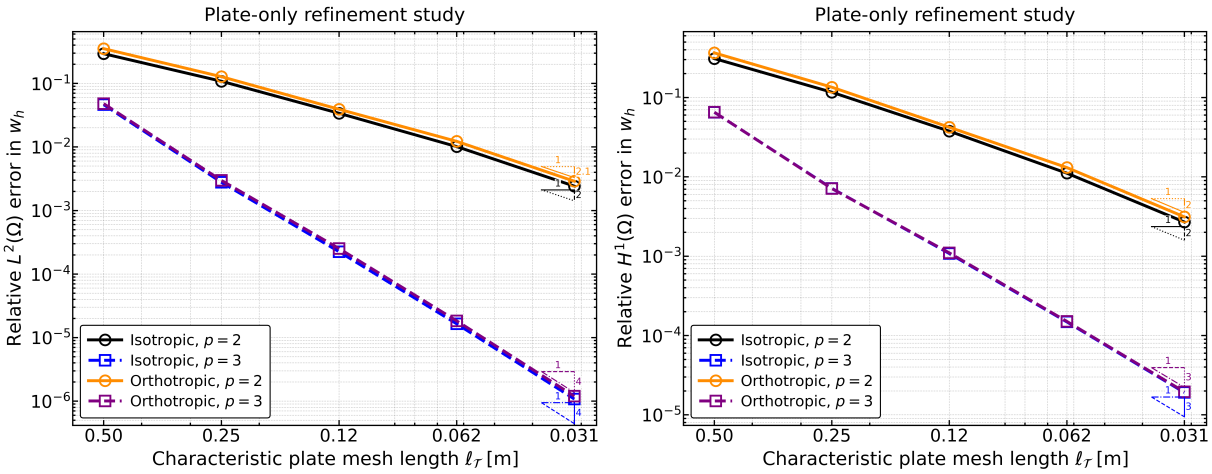
## Plate-only mesh refinement

The plate-only study sets $\theta_h = 0$, $k_s = 0$, and $\rho_m = 0$. It therefore isolates the unsupported plate discretization under

$$q_{\text{mech}} = -10^3 \,\text{N m}^{-2}.$$

The mesh resolution argument is

$$N_\Omega \in \{2, 4, 8, 16, 32\}, \qquad p \in \{2, 3\}, \qquad \ell_{\mathcal{T}} = \frac{1}{N_\Omega} \,\text{m}. \tag{7.29}$$




**(a) Relative $L^2(\Omega)$ error** **(b) Relative $H^1(\Omega)$ error**

**Figure 7.2:** Plate-only refinement for $\theta_h = 0$, $k_s = 0$, and $q_{\mathrm{mech}} = -10^3\ \mathrm{N\,m^{-2}}$. The reference is the refined FOM solution with $N_\Omega = 64$ and $p = 4$. The rate markers show local slopes over the last refinement interval.

The reference is a refined FOM solution with $N_\Omega = 64$ and $p = 4$. The two rigidity sets are

$$\begin{aligned} \text{isotropic:}\quad & D_x = D_y = 10^4\ \mathrm{N\,m}, \qquad D_{xy} = 0.3 \times 10^4\ \mathrm{N\,m}, \\ & D_s = 0.35 \times 10^4\ \mathrm{N\,m}, \\ \text{orthotropic:}\quad & D_x = 2 \times 10^4\ \mathrm{N\,m}, \qquad D_y = 10^4\ \mathrm{N\,m}, \\ & D_{xy} = 0.25 \times 10^4\ \mathrm{N\,m}, \qquad D_s = 0.30 \times 10^4\ \mathrm{N\,m}. \end{aligned}$$

Let $w_h^{\mathrm{fine}}$ denote the refined numerical reference. The dimensionless errors are

$$\varepsilon_{w,\mathrm{rel}}^{L^2} := \frac{\|w_h - w_h^{\mathrm{fine}}\|_{L^2(\Omega)}}{\|w_h^{\mathrm{fine}}\|_{L^2(\Omega)}}, \qquad \varepsilon_{w,\mathrm{rel}}^{H^1} := \frac{\|w_h - w_h^{\mathrm{fine}}\|_{H^1(\Omega)}}{\|w_h^{\mathrm{fine}}\|_{H^1(\Omega)}}. \tag{7.30}$$

**Table 7.1:** Plate-only rates obtained by least-squares fitting of $\log(\varepsilon)$ against $\log(\ell_{\mathcal{T}})$ over $N_\Omega = 8, 16, 32$, with the finest-level relative errors.

| case | $p$ | $r_w^{L^2}$ | $r_w^{H^1}$ | finest $\varepsilon_{w,\mathrm{rel}}^{L^2}$ | finest $\varepsilon_{w,\mathrm{rel}}^{H^1}$ |
|---|---|---|---|---|---|
| isotropic | 2 | 1.90 | 1.90 | $2.42 \times 10^{-3}$ | $2.70 \times 10^{-3}$ |
| isotropic | 3 | 3.85 | 2.91 | $1.08 \times 10^{-6}$ | $1.92 \times 10^{-5}$ |
| orthotropic | 2 | 1.88 | 1.88 | $2.89 \times 10^{-3}$ | $3.13 \times 10^{-3}$ |
| orthotropic | 3 | 3.86 | 2.91 | $1.20 \times 10^{-6}$ | $1.95 \times 10^{-5}$ |

Figure 7.2 shows a monotone decrease of both the relative $L^2(\Omega)$ and $H^1(\Omega)$ errors under mesh refinement. For each polynomial degree, the isotropic and orthotropic curves are nearly parallel, while the $p = 3$ approximation exhibits a substantially steeper error decay than $p = 2$. Table 7.1 quantifies this behavior: for $p = 2$, the fitted rates are 1.88–1.90 in both norms, whereas for $p = 3$ they are 3.85–3.86 in $L^2(\Omega)$ and 2.91 in $H^1(\Omega)$. The close agreement between the isotropic and orthotropic rates indicates that the observed convergence order is essentially unchanged by the rigidity set in this test.

For the symmetric $C^0$ interior-penalty method, the analysis of Wells and

Dung distinguishes the quadratic case from degrees $p \geq 3$ [27]. Under the corresponding regularity assumptions,

$$\|w - w_h\|_{L^2(\Omega)} \leq \begin{cases} C\ell_{\mathcal{T}}^{p}\|w\|_{H^{p+1}(\Omega)}, & p = 2, \\ C\ell_{\mathcal{T}}^{p+1}\|w\|_{H^{p+1}(\Omega)}, & p \geq 3. \end{cases} \tag{7.31}$$

The observed second-order $L^2$ behavior for $p = 2$, rather than a formal third-order rate, is therefore the relevant theoretical comparison. The approximately fourth-order $L^2$ and third-order $H^1$ behavior for $p = 3$ is also consistent with the plate-only expectation.

## Coupled refinement, iteration convergence, and computational cost

The coupled refinement study retains the displacement dependence of $h_c$ and follows the partitioned strategy of Section 7.5. For both material cases,

$$k_s = 10^3\ \mathrm{N\,m^{-3}}, \qquad q_{\mathrm{mech}} = -10^4\ \mathrm{N\,m^{-2}}, \qquad \rho_m = 1050\ \mathrm{kg\,m^{-3}},$$

so self-weight enters through Eq. (6.29). For this controlled refinement test, $\varepsilon_{\mathrm{r}} = 0$ and $\varepsilon_f = 1$. Radiation is therefore disabled and the foundation is bilateralized, while the displacement dependence of the contact–gap exchange law remains active. The study assesses coupled spatial refinement and convergence of the partitioned iteration for the stated parameter configuration; it does not test the radiative or unilateral nonlinearities.

The rigidity values are those used in the plate-only isotropic and orthotropic cases. The thermal data are

$$\begin{aligned} T_{\mathrm{amb}} &= 300.15\ \mathrm{K}, & T_{\mathrm{sub}} &= 297.15\ \mathrm{K}, \\ h_{\mathrm{con}} &= 10\ \mathrm{W\,m^{-2}K^{-1}}, & q_s &= 700\ \mathrm{W\,m^{-2}}, \\ h_c^{\mathrm{cont}} &= 200\ \mathrm{W\,m^{-2}K^{-1}}, & h_c^{\mathrm{gap}} &= 5\ \mathrm{W\,m^{-2}K^{-1}}, \\ \eta_c &= 50\ \mathrm{m^{-1}}, & k_x = k_y = k_z &= 0.35\ \mathrm{W\,m^{-1}K^{-1}}, \\ \alpha_x = \alpha_y &= 1.3 \times 10^{-4}\ \mathrm{K^{-1}}. \end{aligned}$$

At each iteration, the current displacement updates the lower-face exchange coefficient, the three-dimensional temperature field is recomputed, and the reduced thermal driver enters the plate solve.

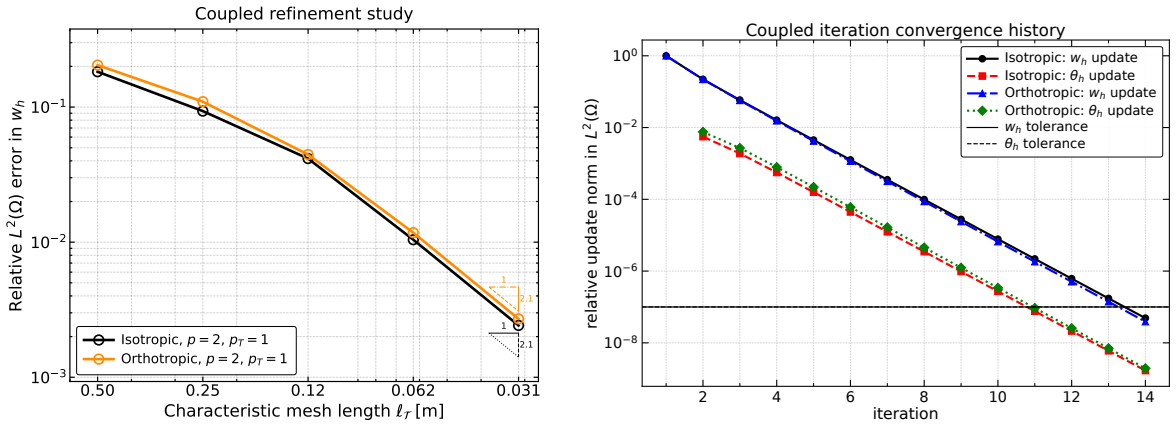
At level $N_\Omega$, the plate mesh is the triangular mesh generated with resolution parameter $N_\Omega$, while the heat mesh contains

$$N_\Omega \times (N_\Omega/2) \times \max(1, N_\Omega/4)$$

box cells before tetrahedral subdivision. The nominal plate length is $\ell_{\mathcal{T}} = 1/N_\Omega$ m, with

$$N_\Omega \in \{2, 4, 8, 16, 32\}, \qquad (p, p_T, p_\theta) = (2, 1, 2) \tag{7.32}$$

for both the isotropic and orthotropic rigidity sets. The through-thickness reduction uses 8 quadrature points for $N_\Omega = 2, 4, 8, 16$ and 10 points for $N_\Omega = 32$. The converged refined reference uses $N_\Omega = 64$, $p = 4$, a $64 \times 32 \times 16$ heat mesh with $p_T = 2$, $p_\theta = 4$, and 24 through-thickness quadrature points.




**(a) Relative displacement error** **(b) Refined-reference update histories**

**Figure 7.3:** Coupled refinement for the isotropic and orthotropic cases with $(p, p_T, p_\theta) = (2, 1, 2)$. Panel (a) shows the relative $L^2(\Omega)$ displacement errors and final-interval slope markers. Panel (b) shows the refined-reference updates of $w_h$ and $\theta_h$, with the common tolerance $10^{-7}$. Fitted displacement rates are listed in Table 7.2.

**Table 7.2:** Least-squares log–log coupled displacement rates over $N_\Omega = 8, 16, 32$, with relative displacement and thermal-driver errors at $N_\Omega = 32$ against the refined coupled FOM references.

| case | $(p, p_T, p_\theta)$ | $r_w^{L^2}$ | finest $\varepsilon_{w,\mathrm{rel}}^{L^2}$ | finest $\varepsilon_{\theta,\mathrm{rel}}^{L^2}$ |
|---|---|---|---|---|
| isotropic | $(2, 1, 2)$ | 2.05 | $2.42 \times 10^{-3}$ | $1.14 \times 10^{-4}$ |
| orthotropic | $(2, 1, 2)$ | 2.02 | $2.71 \times 10^{-3}$ | $1.45 \times 10^{-4}$ |

In addition to the displacement error in Eq. (7.30), with $w_h^{\mathrm{fine}}$ now denoting the converged coupled reference displacement, the thermal-driver error is

$$\varepsilon_{\theta,\mathrm{rel}}^{L^2} := \frac{\|\theta_h - \theta_h^{\mathrm{fine}}\|_{L^2(\Omega)}}{\|\theta_h^{\mathrm{fine}}\|_{L^2(\Omega)}}. \tag{7.33}$$

Here $\theta_h^{\mathrm{fine}}$ is the thermal driver from the same converged coupled reference state. The partitioned parameters are

$$\omega = 0.7, \qquad \mathrm{tol}_w = \mathrm{tol}_\theta = 10^{-7}, \qquad K_{\max} = 50. \tag{7.34}$$

Figure 7.3(a) shows that the relative displacement error decreases monotonically for both rigidity sets. Table 7.2 reports fitted rates of 2.05 and 2.02 for the isotropic and orthotropic cases, respectively, together with finest-level displacement errors of $2.42 \times 10^{-3}$ and $2.71 \times 10^{-3}$ and thermal-driver errors of $1.14 \times 10^{-4}$ and $1.45 \times 10^{-4}$. The approximately second-order displacement rates are consistent with the behavior relevant to the quadratic $C^0$ interior-penalty plate approximation. These coupled rates also reflect the three-dimensional thermal approximation, the thermal-to-plate reduction, transfer between nonmatching spaces, and comparison with a refined coupled reference.

Testing nominal fourth-order coupled $L^2$ displacement behavior under suitable regularity and asymptotic assumptions would require at least the fully cubic choice $(p, p_T, p_\theta) = (3, 3, 3)$, together with a sufficiently resolved higher-order coupled reference. Because every partitioned iteration would then require a cubic three-dimensional heat solve, such a campaign is computationally too expensive for the present resources. The fully cubic case is therefore not computed, and no numerical fourth-order claim is made for the coupled problem.

**Iteration counts and measured computational cost.** Each partitioned iteration comprises one three-dimensional thermal solve, the through-thickness reduction and transfer to $\Theta_h$, one two-dimensional mechanical solve, displacement relaxation, and two relative-update evaluations. On the refined reference discretization, the thermal system has 276,705 degrees of freedom, whereas the mechanical and thermal-driver spaces each have 81,469 degrees of freedom. Considering only the two PDE systems, the thermal unknowns account for 77.3% and the mechanical unknowns for 22.7% of their combined algebraic size; the thermal system is 3.40 times larger. These percentages are algebraic-size proxies and are reported separately from the measured wall-clock split.

The retained convergence data consist of ten refinement runs, obtained from two material cases over five mesh levels, together with two refined reference runs. All twelve runs converge after 14 partitioned iterations. Each run therefore requires 14 thermal and 14 mechanical PDE solves, or 28 PDE solves per run, in addition to the reductions, relaxations, and stopping checks. Figure 7.3(b) shows that the $\theta_h$ update first falls below $10^{-7}$ at iteration 11, whereas the displacement update reaches the same tolerance at iteration 14 and determines termination.

For the wall-clock assessment, the two refined reference cases shown in Figure 7.3(b) were rerun once each with the same discretization and solver parameters. Timing begins immediately before the complete partitioned solve and ends after the final coupled state is returned. Module import and plate-domain setup are excluded. Thermal timing includes heat-mesh and form construction, assembly, the PETSc solve, and construction of the temperature increment. Mechanical timing includes plate-form construction, assembly, the PETSc solve, and reconstruction of the global displacement field. The first iteration contains form compilation and cache initialization; it is included in the total wall time but omitted from the representative per-iteration statistics.

From iterations 2–14, the median wall-clock cost per iteration is 230.89 s for the isotropic reference and 364.24 s for the orthotropic reference. Across the 26 retained iteration samples, the median is 265.21 s, with an interquartile range from 231.85 to 357.97 s. Over these iterations, the three-dimensional thermal PDE stage accounts for 85.4% of the measured wall time, the thermal-to-plate reduction and transfer for 10.2%, the mechanical PDE stage for 4.4%, and relaxation, convergence checks, state copies, and the remaining coupling operations for less than 0.1%. Grouping the reduction and transfer with the thermal subproblem gives measured thermal and mechanical shares of 95.6% and 4.4%, respectively. Including the first iteration and finalization, the complete coupled wall times are 3199.68 s (53.3 min) for the isotropic case and 4978.89 s (83.0 min) for the orthotropic case. Since each material case was timed once, these values characterize the stated execution platform and protocol rather than a hardware-independent complexity law.

## Scope of the verification evidence

The verification evidence supports the discrete ingredients exercised by the reported tests [35, 36]. The Navier comparison checks the plate field for a single-panel, simply supported, constant-foundation problem with a fixed reconstructed thermal driver; it is neither a pointwise unilateral solution nor a solution of the displacement-dependent fixed

point. The plate-only refinement isolates the symmetric $C^0$ interior-penalty discretization. The coupled refinement verifies simultaneous thermal and mechanical refinement and fixed-point convergence for the controlled $(2, 1, 2)$ sequences, with radiation disabled and the foundation made bilateral. The panel-interface discretization is inherited from, and assessed in, Chapter 4; it is not revalidated here through a separate thermomechanical partition-layout sweep.

Appendix A records the implementation architecture, software stack, data-transfer operations, and timing conventions required to reproduce the studies.

## 7.8 Summary of the thermomechanical full-order model

The thermomechanical full-order realization combines a continuous Lagrange discretization of the three-dimensional thermal problem with a panel-wise $C^0$ plate discretization. The thermal mesh resolves heat conduction and the nonlinear radiative boundary term, while numerical through-thickness integration and interpolation reconstruct the plate thermal driver. Symmetric interior-penalty stabilization acts on numerical facets within each physical panel, whereas separate displacement and normal-rotation penalties transmit the mechanical response across physical panel interfaces.

The auxiliary Navier comparison provides a controlled analytical benchmark for a simplified constant-coefficient problem. The plate-only refinement reproduces the expected distinction between quadratic and cubic $C^0$ interior-penalty behavior. The coupled $(2, 1, 2)$ sequences show monotone, approximately second-order displacement convergence for both rigidity sets, and all twelve retained coupled runs terminate after 14 iterations. The measured refined-reference profile further shows that the three-dimensional thermal stage and the thermal-to-plate reduction account for 95.6% of the recurrent iteration time on the stated platform. A fully cubic coupled refinement is not reported because its three-dimensional thermal cost is prohibitive for the available resources.

The accepted full-order states comprise the global displacement, thermal driver, and associated convergence metadata. They provide the common source data for the two one-field workflows developed in Chapter 8.

# Reduced-Order Modeling of the Thermomechanical Plate Problem

# 8


8.1 Thermomechanical many-query setting and reduction objective . . . . . 87
8.2 Reduced output maps and snapshot databases . . . . 87
8.3 Metric-aware field-wise POD spaces . . . . . 89
8.4 Hybrid projection-based thermomechanical ROM 90
8.5 Non-intrusive one-field ROMs . . . . . 91
8.6 Error measures and assessment protocol . . . . . 93
8.7 Numerical assessment . . 94
8.8 Interpretation, limitations, and chapter conclusions . 100


## 8.1 Thermomechanical many-query setting and reduction objective

Chapters 6 and 7 established the physically coupled thermomechanical model and its high-fidelity finite element realization. For every parameter value, the accepted full-order state is obtained only after the partitioned thermal–mechanical iteration has satisfied the displacement and thermal-driver stopping criteria. Repeating that calculation over a parameter domain requires several three-dimensional thermal and plate-mechanical solves and motivates the reduced descriptions developed in this chapter. Related model-reduction studies address parameterized three-dimensional heat transfer, steady one-way thermomechanics, and thermomechanical systems reduced through bases assembled from separate physical fields [44, 45, 91].

The reduction is deliberately output-specific. Two complementary one-field quantities are considered: the global transverse displacement $\widehat{w}_h(\boldsymbol{\mu})$ and the global plate thermal driver $\widehat{\theta}_h(\boldsymbol{\mu})$. The qualifiers "displacement-only" and "thermal-driver-only" describe the selected ROM output; they do not describe the physics used to generate the snapshots. Both fields originate from the same accepted, converged, two-way coupled FOM state.

Two computational strategies are examined. The first is a hybrid mechanical POD–Galerkin formulation that projects only the displacement correction while retaining the three-dimensional thermal solve, the thermal-to-plate reduction, and full-order mechanical residual and Jacobian assembly. The second comprises non-intrusive one-field surrogates for either displacement or thermal driver: POD interpolation, Gaussian-process and neural-network prediction of POD coordinates, POD-enhanced deep learning, and direct full-state deep learning. The hybrid construction preserves the partitioned FOM iteration and is therefore not a fully reduced thermomechanical solver. In contrast, a non-intrusive model replaces the complete online partitioned solve only for its selected one-field output.

Training data are used to construct the POD bases and fit the reduced models. Validation data are used to choose reduced or latent dimensions and to select the best neural-network checkpoint. Test data are reserved for the final accuracy, timing, and spatial-error assessment. Direct POD projection is included only to measure the error introduced by truncating the basis; it requires the corresponding full-order field and is therefore not a predictive ROM.

## 8.2 Reduced output maps and snapshot databases

Three snapshot campaigns are considered. In each campaign, $\boldsymbol{\mu} \in \mathcal{D} \subset \mathbb{R}^P$ denotes the active parameter vector and $P = \dim(\boldsymbol{\mu})$ its dimension. The

campaign labels are dimension-based rather than sequential: Case 2 has two active parameters, Case 3 has three, and Case 5 has five. Table 8.1 lists the corresponding parameter domains.

| case $P$ | configuration | active parameters and ranges |
|---|---|---|
| 2 | monolithic | $q_{\mathrm{mech}} \in [-6500, -1500]\ \mathrm{N\,m^{-2}}$ |
| | $\boldsymbol{\mu} = (q_{\mathrm{mech}}, q_s)$ | $q_s \in [150, 850]\ \mathrm{W\,m^{-2}}$ |
| 3 | two-panel contact | $k_s \in [10^5, 2\times10^6]\ \mathrm{N\,m^{-3}}$ |
| | $\boldsymbol{\mu} = (k_s, T_{\mathrm{amb}}, T_{\mathrm{sub}})$ | $T_{\mathrm{amb}} \in [298.15, 313.15]\ \mathrm{K}$ |
| | | $T_{\mathrm{sub}} \in [295.15, 310.15]\ \mathrm{K}$ |
| 5 | monolithic | $k_s \in [10^5, 3\times10^6]\ \mathrm{N\,m^{-3}}$ |
| | $\boldsymbol{\mu} = (k_s, q_{\mathrm{mech}}, q_s,$ | $q_{\mathrm{mech}} \in [-9000, -2000]\ \mathrm{N\,m^{-2}}$ |
| | $h_c^{\mathrm{cont}}, h_c^{\mathrm{gap}})$ | $q_s \in [150, 850]\ \mathrm{W\,m^{-2}}$ |
| | | $h_c^{\mathrm{cont}} \in [75, 300]\ \mathrm{W\,m^{-2}K^{-1}}$ |
| | | $h_c^{\mathrm{gap}} \in [3, 15]\ \mathrm{W\,m^{-2}K^{-1}}$ |

**Table 8.1:** Active parameter domains for the three ROM campaigns. The case label equals $P = \dim(\boldsymbol{\mu})$, the number of active parameters.

Table 8.1 shows that Cases 2 and 5 are monolithic, whereas Case 3 contains two panels separated by one vertical physical interface. Case 2 varies mechanical loading and absorbed solar flux; Case 3 varies foundation stiffness and the two boundary temperatures; and Case 5 varies foundation stiffness, mechanical loading, absorbed solar flux, and the contact- and gap-side heat-transfer coefficients. All three configurations use a free exterior edge and a localized patch load.

For each $\boldsymbol{\mu}$, the converged partitioned computation produces the accepted displacement $w_h(\boldsymbol{\mu})$, the thermal driver $\theta_h(\boldsymbol{\mu})$, and the associated three-dimensional thermal field. The reduced outputs are stored in global continuous finite element spaces $X_h^w$ and $X_h^\theta$. The corresponding transfer maps are

$$\begin{aligned} \mathcal{P}_h^w &: V_h^w \longrightarrow X_h^w, \quad \widehat{w}_h(\boldsymbol{\mu}) \coloneqq \mathcal{P}_h^w w_h(\boldsymbol{\mu}), \\ \mathcal{P}_h^\theta &: V_h^\theta \longrightarrow X_h^\theta, \quad \widehat{\theta}_h(\boldsymbol{\mu}) \coloneqq \mathcal{P}_h^\theta \theta_h(\boldsymbol{\mu}). \end{aligned} \tag{8.1}$$

For the monolithic configurations, these maps reduce to the finite element data transfer used to construct the snapshot arrays. In the two-panel contact configuration, $\mathcal{P}_h^w$ transfers the panel-wise mechanical solution to the global continuous representation used by the ROM. The thermal driver is already stored on the global plate mesh; its Case 3 output space has a slightly smaller dimension because of the interface-conforming mesh topology. Consequently, the Case 3 errors reported below quantify global fields; panel-interface jumps are not introduced as additional reduced outputs.

The two high-fidelity output maps are

$$\begin{aligned} \mathcal{S}_h^w &: \mathcal{D} \longrightarrow X_h^w, \quad \mathcal{S}_h^w(\boldsymbol{\mu}) = \widehat{w}_h(\boldsymbol{\mu}), \\ \mathcal{S}_h^\theta &: \mathcal{D} \longrightarrow X_h^\theta, \quad \mathcal{S}_h^\theta(\boldsymbol{\mu}) = \widehat{\theta}_h(\boldsymbol{\mu}). \end{aligned} \tag{8.2}$$

Both maps are constructed from the same case-specific thermomechanical snapshot campaign. Neither output ROM reduces the complete three-dimensional temperature field.

All snapshot campaigns use the same high-fidelity discretization:

$$N_\Omega = 64, \quad p = 2, \qquad (N_x^T, N_y^T, N_z^T) = (64, 32, 16), \quad p_T = 1, \quad p_\theta = 1. \tag{8.3}$$

The through-thickness reduction is evaluated with 20 quadrature points. The partitioned coupling settings are

$$\omega = 0.7, \qquad \mathrm{tol}_w = \mathrm{tol}_\theta = 10^{-5}, \qquad K_{\max} = 25. \tag{8.4}$$

For a reduced field $\chi \in \{w, \theta\}$, the stored database is

$$\mathcal{E}_h^\chi = \left\{ \left(\boldsymbol{\mu}^{(j)}, \widehat{\chi}_h^{(j)}\right) \right\}_{j=1}^{N_s}, \qquad \widehat{\chi}_h^{(j)} = \mathcal{S}_h^\chi(\boldsymbol{\mu}^{(j)}). \tag{8.5}$$

Only accepted, converged full-order states are retained in the ROM databases. The final databases contain $N_s = 450$, 378, and 500 records for Cases 2, 3, and 5, respectively.

Within each case, the accepted records are partitioned once into training, validation, and test subsets in the proportions 80%, 10%, and 10%. Seed 100 determines the training subset, and seed 101 divides the remaining records equally between validation and test samples. The same case-specific indices are used for the displacement-only and thermal-driver-only workflows, so each sample belongs to the same statistical subset in both output analyses [54].

Table 8.2 summarizes the accepted database sizes, sample partitions, and global output-space dimensions. The two monolithic campaigns use the same global output-space dimensions, whereas the interface-conforming Case 3 mesh gives slightly smaller displacement and thermal-driver spaces. These dimensions refer to the full-order output spaces; the reduced dimensions are selected later from the POD and validation studies.

| case $P$ | $N_s$ | train/validation/test | dim $X_h^w$ | dim $X_h^\theta$ |
|---|---|---|---|---|
| 2 | 450 | 360/45/45 | 20,503 | 5,194 |
| 3 | 378 | 302/38/38 | 20,375 | 5,160 |
| 5 | 500 | 400/50/50 | 20,503 | 5,194 |

**Table 8.2:** Accepted snapshot counts, train/validation/test partitions, and global output-space dimensions. The case label equals $P = \dim(\boldsymbol{\mu})$, the number of active parameters.

## 8.3 Metric-aware field-wise POD spaces

Proper orthogonal decomposition is applied to the raw, uncentered snapshots in the training subset [40, 41, 53]. Because displacement and thermal-driver fields have different physical dimensions and different spatial regularity requirements, they are reduced in separate metrics. The displacement POD uses the $H^1$-type metric adopted for the mechanical field,

$$(u_h, v_h)_w := \int_\Omega u_h v_h \, \mathrm{d}\Omega + \int_\Omega \nabla u_h \cdot \nabla v_h \, \mathrm{d}\Omega, \qquad \|v_h\|_w^2 = (v_h, v_h)_w, \tag{8.6}$$

whereas the thermal-driver POD uses the $L^2(\Omega)$ metric

$$(\eta_h, \xi_h)_\theta := \int_\Omega \eta_h \xi_h \, \mathrm{d}\Omega, \qquad \|\xi_h\|_\theta^2 = (\xi_h, \xi_h)_\theta. \tag{8.7}$$

These are numerical reduction metrics. In particular, Eq. (8.6) is not identified with the complete thermomechanical energy.

For $\chi \in \{w, \theta\}$, let $(\cdot,\cdot)_\chi$ denote the corresponding metric and let $N_{\mathrm{tr}}$ be the number of training rows in the case under consideration. The training

correlation matrix is

$$C_{ij}^{\chi} = \frac{1}{N_{\mathrm{tr}}} \left( \widehat{\chi}_h^{(i)}, \widehat{\chi}_h^{(j)} \right)_{\chi}, \qquad i, j = 1, \ldots, N_{\mathrm{tr}}. \tag{8.8}$$

If $(\lambda_n^{\chi}, \boldsymbol{v}_n^{\chi})$ is an eigenpair of $C^{\chi}$, the corresponding normalized POD mode is

$$\zeta_n^{\chi} = \frac{1}{\sqrt{N_{\mathrm{tr}} \lambda_n^{\chi}}} \sum_{j=1}^{N_{\mathrm{tr}}} (\boldsymbol{v}_n^{\chi})_j \widehat{\chi}_h^{(j)}, \qquad (\zeta_n^{\chi}, \zeta_m^{\chi})_{\chi} = \delta_{nm}. \tag{8.9}$$

The retained field space and cumulative fraction of the computed 30-mode POD spectrum are

$$X_{N_{\chi}}^{\chi} \coloneqq \operatorname{span}\{\zeta_1^{\chi}, \ldots, \zeta_{N_{\chi}}^{\chi}\}, \qquad \mathscr{E}_{N_{\chi}}^{\chi} \coloneqq \frac{\sum_{n=1}^{N_{\chi}} \lambda_n^{\chi}}{\sum_{n=1}^{N_{\max}} \lambda_n^{\chi}}, \qquad N_{\max} = 30. \tag{8.10}$$

Accordingly, $\mathscr{E}_{N_{\chi}}^{\chi}$ denotes a cumulative fraction of the first 30 computed POD eigenvalues.

The POD computation uses a numerical tolerance of $10^{-15}$. After the basis is fixed, any validation or test field $v_h$ can be projected according to

$$\Pi_{N_{\chi}}^{\chi} v_h = \sum_{n=1}^{N_{\chi}} a_n^{\chi}(v_h) \zeta_n^{\chi}, \qquad a_n^{\chi}(v_h) = (v_h, \zeta_n^{\chi})_{\chi}. \tag{8.11}$$

Projection requires the high-fidelity field and is therefore used only to quantify basis-truncation error; it is a non-predictive reference, not an online ROM.

## 8.4 Hybrid projection-based thermomechanical ROM

The projection-based model acts on the plate displacement while retaining the high-fidelity thermal subproblem. Let $Z_N^w \in \mathbb{R}^{N_h^w \times N}$ contain the finite element coefficient vectors of the displacement POD modes and let $\boldsymbol{a} \in \mathbb{R}^N$ denote the reduced coordinates. In a panel-partitioned configuration, $N_h^w$ refers to the corresponding broken mechanical space, so the foundation, interface, and contact contributions remain part of the assembled high-fidelity operators. During a mechanical step, the displacement is approximated by

$$w_h \approx Z_N^w \boldsymbol{a}. \tag{8.12}$$

Writing the discrete mechanical residual from Eq. (7.17) in coefficient form as $\boldsymbol{F}_h(w_h; \theta_h, \boldsymbol{\mu})$, the reduced equilibrium condition is

$$\boldsymbol{r}_N(\boldsymbol{a}; \theta_h, \boldsymbol{\mu}) \coloneqq (Z_N^w)^{\top} \boldsymbol{F}_h(Z_N^w \boldsymbol{a}; \theta_h, \boldsymbol{\mu}) = \boldsymbol{0}. \tag{8.13}$$

A reduced Newton step uses

$$J_N = (Z_N^w)^{\top} J_h(Z_N^w \boldsymbol{a}; \theta_h, \boldsymbol{\mu}) Z_N^w, \qquad J_N \, \delta \boldsymbol{a} = -\boldsymbol{r}_N, \tag{8.14}$$

where $J_h$ is the full-order plate Jacobian.

The thermomechanical loop follows the same ordering, relaxation, and stopping criteria as the FOM. At every partitioned iteration, however, the three-dimensional heat problem is solved in $V_{T,h}$, the reduction to $\theta_h$ is carried out in the full plate thermal space, and the full-order mechanical residual and Jacobian are assembled before projection. Structure-preserving reduction of coupled systems provides broader methodological context for retaining the architecture of the parent coupled problem [92]. This formulation is therefore termed a hybrid thermomechanical POD–Galerkin model: it reduces the mechanical linear algebra but does not eliminate the high-dimensional thermal solve or the full-order assembly costs.

The projected residual and Jacobian definitions retain the complete mechanical operator before projection. Hence, the unilateral foundation, interior-penalty, and physical-interface contributions require no separate reduced constitutive approximation. The formulation preserves the original partitioned update and stopping criteria, while the mechanical correction is computed in the POD space.

## 8.5 Non-intrusive one-field ROMs

For a single field $\chi \in \{w, \theta\}$, the online problem is reduced to approximating the coefficient map

$$\boldsymbol{a}^{\chi} : \mathcal{D} \longrightarrow \mathbb{R}^{N_\chi}, \qquad \boldsymbol{a}^{\chi}(\boldsymbol{\mu}) = \left(a_1^{\chi}(\widehat{\chi}_h(\boldsymbol{\mu})), \ldots, a_{N_\chi}^{\chi}(\widehat{\chi}_h(\boldsymbol{\mu}))\right)^{\top}, \tag{8.15}$$

followed by the reconstruction

$$\widehat{\chi}_{N_\chi}(\boldsymbol{\mu}) = \sum_{n=1}^{N_\chi} \widehat{a}_n^{\chi}(\boldsymbol{\mu})\zeta_n^{\chi}. \tag{8.16}$$

The active parameters and POD coefficients are standardized componentwise using statistics from the training subset.

**POD interpolation.** PODI-RBF reconstructs the coefficient map by vector-valued radial-basis interpolation with a thin-plate-spline kernel and zero smoothing [81, 82]. PODI-Linear uses simplicial piecewise-linear interpolation in the normalized parameter space. When a query lies outside the convex hull of the training points and the linear interpolant is undefined, a nearest-neighbor coefficient vector is used as a documented fallback.

**POD Gaussian-process regression.** POD-GPR fits one independent Gaussian process to each standardized POD coefficient [48, 77, 93]. The covariance model combines a constant kernel, an anisotropic squared-exponential kernel, and a white-noise contribution. A diagonal jitter of $10^{-9}$ and multiple optimizer restarts are used. The displacement and thermal-driver models contain $N_w$ and $N_\theta$ scalar regressors, respectively. Posterior means supply the predicted coefficients. Posterior standard deviations are retained by the implementation but are not propagated to probabilistic field bands in the present assessment.

**POD neural network.** POD–NN maps the normalized parameters directly to the standardized POD coordinate vector through a fully connected network with hidden widths (128, 96, 64), ELU activations, and a linear output [47, 60]. The network is trained directly on the precomputed standardized POD projection coordinates by minimizing the coefficient mean-square error. AdamW optimization is combined with validation-based learning-rate reduction, early stopping, and restoration of the minimum-validation-loss checkpoint. At prediction time, the coordinates are inverse-scaled and lifted through the fixed POD basis according to Equation (8.16).

**POD-enhanced deep-learning ROM.** The POD-DL-ROM introduces a nonlinear latent representation of the POD coefficient vector. Let $E_c : \mathbb{R}^{N_\chi} \to \mathbb{R}^{n_z}$ and $D_c : \mathbb{R}^{n_z} \to \mathbb{R}^{N_\chi}$ denote the coefficient encoder and decoder, and let $G_\vartheta : \mathbb{R}^P \to \mathbb{R}^{n_z}$ be the parameter-to-latent network. The predicted coefficient vector is

$$\widehat{\boldsymbol{a}}_{\text{POD-DL}}(\boldsymbol{\mu}) = D_c\left(G_\vartheta(\boldsymbol{\mu})\right), \tag{8.17}$$

which is lifted by Equation (8.16). The encoder, decoder, and latent map use ELU multilayer perceptrons, with reversed hidden widths in the decoder. Training combines coefficient autoencoding, parameter-to-latent regression, and joint latent-map/decoder refinement. This construction is consistent with POD-enhanced deep-learning reductions in which a linear high-dimensional compression is followed by a learned nonlinear manifold [49].

**Direct deep-learning ROM.** DL-ROM bypasses the POD coefficients and learns a nonlinear encoding of the full output vector [50]. The displacement encoder uses four hidden layers with widths 2048, 1024, 512, and 256, while the parameter-to-latent map uses widths (128, 96, 64). The thermal-driver model uses the analogous field-specific construction. Each model is trained and evaluated independently for its physical output.

The decoder-level objectives used by POD–DL-ROM and direct DL-ROM are evaluated on normalized finite-dimensional output vectors. These training objectives are distinct from the finite element mass-matrix error used for assessment and from the POD metrics in Equations (8.6) and (8.7).

Table 8.3 summarizes the learned maps and principal implementation settings of the six non-intrusive one-field methods. The first four methods approximate the parameter-to-POD-coordinate map directly, whereas POD-DL-ROM introduces a nonlinear latent representation of the POD coordinates and direct DL-ROM reconstructs the selected full output field through a field-specific decoder.

The exact layer widths, latent dimensions, encoder–decoder paths, and deployed online paths for POD–NN, POD–DL-ROM, and direct DL-ROM are centralized in Section A.5. The displacement and plate-thermal-driver models remain independent throughout preprocessing, training, checkpointing, online evaluation, and field reconstruction.

**Table 8.3:** Non-intrusive one-field ROM maps and principal settings.

| method | learned map | principal setting |
|---|---|---|
| PODI-RBF | $\boldsymbol{\mu} \mapsto \boldsymbol{a}$ | thin-plate-spline RBF, zero smoothing |
| PODI-Linear | $\boldsymbol{\mu} \mapsto \boldsymbol{a}$ | simplicial interpolation, nearest-neighbor fallback |
| POD-GPR | $\boldsymbol{\mu} \mapsto a_n$, independently | constant × anisotropic RBF + white-noise kernel |
| POD–NN | $\boldsymbol{\mu} \mapsto \boldsymbol{a}$ | ELU MLP (128, 96, 64), coefficient MSE, validation checkpointing |
| POD-DL-ROM | $\boldsymbol{\mu} \mapsto \boldsymbol{z} \mapsto \boldsymbol{a}$ | coefficient autoencoder, latent map, and decoder refinement |
| DL-ROM | $\boldsymbol{\mu} \mapsto \boldsymbol{z} \mapsto \widehat{\chi}_h$ | field-specific full-state autoencoder |

## 8.6 Error measures and assessment protocol

All spatial errors reported for the validation and test subsets use the field-appropriate finite element mass matrix $M_\chi$. For $\chi \in \{w, \theta\}$, the relative $L^2(\Omega)$ error of sample $j$ is

$$e^{(j)}_{L^2,\chi} := \frac{\sqrt{(\widehat{\boldsymbol{\chi}}^{(j)}_h - \widehat{\boldsymbol{\chi}}^{(j)}_{\mathrm{ROM}})^\top M_\chi (\widehat{\boldsymbol{\chi}}^{(j)}_h - \widehat{\boldsymbol{\chi}}^{(j)}_{\mathrm{ROM}})}}{\sqrt{(\widehat{\boldsymbol{\chi}}^{(j)}_h)^\top M_\chi \widehat{\boldsymbol{\chi}}^{(j)}_h}}. \tag{8.18}$$

For the contact configuration, both the FOM and ROM displacement fields are first represented in the global continuous space $X^w_h$, so Eq. (8.18) is evaluated in the same output space for all three cases.

For the spatial diagnostics, the pointwise difference and its absolute value are defined by

$$\Delta\chi^{(j)}_h(\boldsymbol{x}) := \widehat{\chi}^{(j)}_h(\boldsymbol{x}) - \widehat{\chi}^{(j)}_{\mathrm{ROM}}(\boldsymbol{x}), \qquad \varepsilon^{(j)}_{\mathrm{abs},\chi}(\boldsymbol{x}) := \left|\Delta\chi^{(j)}_h(\boldsymbol{x})\right|, \qquad \boldsymbol{x} \in \Omega. \tag{8.19}$$

Thus, $\varepsilon_{\mathrm{abs},w}$ has units of length and $\varepsilon_{\mathrm{abs},\theta}$ has units of temperature per unit length.

For every method $m$, output field $\chi$, and candidate reduced or latent dimension $N$, the validation sweep reports

$$\begin{aligned} \bar{e}^{\mathrm{val}}_{\chi,m,N} &= \frac{1}{N_{\mathrm{val}}} \sum_{j \in \mathcal{F}_{\mathrm{val}}} e^{(j)}_{L^2,\chi,m,N}, \\ e^{\mathrm{val,min}}_{\chi,m,N} &= \min_{j \in \mathcal{F}_{\mathrm{val}}} e^{(j)}_{L^2,\chi,m,N}, \\ e^{\mathrm{val,max}}_{\chi,m,N} &= \max_{j \in \mathcal{F}_{\mathrm{val}}} e^{(j)}_{L^2,\chi,m,N}. \end{aligned} \tag{8.20}$$

The plotted line is the validation mean and the shaded region is the observed samplewise minimum–maximum range over the same validation rows; it is not a statistical confidence interval. These studies assess sensitivity to the linear reduced dimension or nonlinear latent dimension and are used to select the working dimensions for the final cross-case comparison. The test data do not enter this selection.

For each predictive non-intrusive method $m$, reduced output $\chi \in \{w, \theta\}$, and held-out test sample $j$, the matched single-query online speedup is distinguished from the non-predictive POD-projection timing ratio

through

$$S^{(j)}_{\chi,m} := \frac{t^{(j)}_{\mathrm{FOM}}}{t^{(j)}_{\mathrm{ROM},\chi,m}}, \qquad R^{(j)}_{\chi,\mathrm{proj}} := \frac{t^{(j)}_{\mathrm{FOM}}}{t^{(j)}_{\mathrm{proj},\chi}}. \tag{8.21}$$

Here, $t^{(j)}_{\mathrm{FOM}}$ is the wall time of the complete coupled thermomechanical FOM evaluation at $\mu^{(j)}$. Because displacement and thermal driver are extracted from the same accepted coupled state, this full-order timing provides the common numerator for both one-field assessments. The predictive-ROM time $t^{(j)}_{\mathrm{ROM},\chi,m}$ includes parameter preprocessing, evaluation of the trained map, inverse scaling, and reconstruction of the selected output field. It excludes snapshot generation, POD construction, model training, checkpoint loading, and figure generation.

The POD-Proj denominator $t^{(j)}_{\mathrm{proj},\chi}$ measures the evaluation of the projection diagnostic from an already available high-fidelity field. Consequently, $R^{(j)}_{\chi,\mathrm{proj}}$ is reported only as a non-predictive projection-cost diagnostic; it is not interpreted as a deployable online speedup.

## 8.7 Numerical assessment

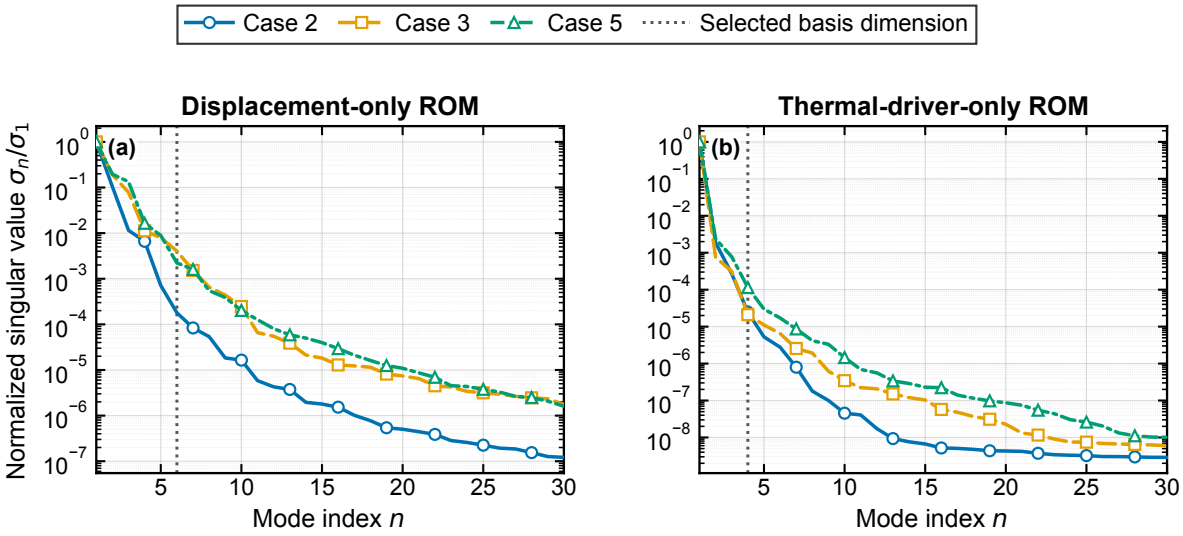



**Figure 8.1:** Training-only normalized POD singular-value decay for the global displacement and thermal-driver output spaces in Cases 2, 3, and 5. The dotted lines identify the working dimensions used in the final comparison.

**POD spectral decay.** Figure 8.1 shows the training-only normalized POD singular-value decay for the displacement and thermal-driver outputs in Cases 2, 3, and 5. Only the spectra of the global output spaces used for non-intrusive prediction and field-error assessment are reported.

Figure 8.1 shows rapid spectral decay for both outputs, with a steeper decrease for the thermal driver. This indicates that fewer thermal-driver modes are needed to reproduce a given fraction of the training snapshots.

The spectra characterize the linear representation capacity of the POD spaces; they do not, by themselves, determine the dimensions used by the predictive ROMs. Those dimensions are selected from the validation study below.

**Validation-based reduced-dimension assessment.** Figure 8.2 presents the two Case 2 validation sweeps, with candidate dimensions

$$N_w \in \{1, 3, 6, 9, 12\}, \qquad N_\theta \in \{1, 4, 8, 11, 15\}.$$

For DL-ROM, the horizontal coordinate denotes the nonlinear latent dimension. Each curve is the mean over the 45 Case 2 validation samples, and each shaded band spans the corresponding samplewise minimum–maximum range.

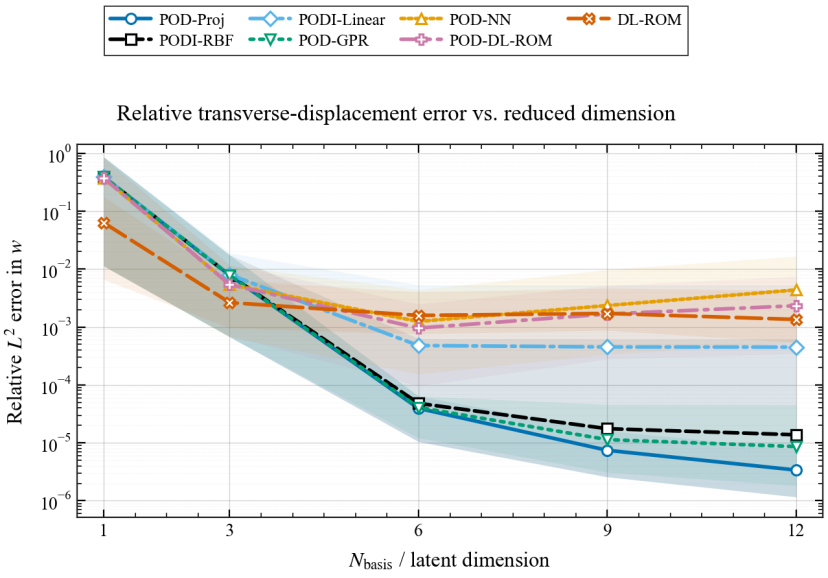
Figure 8.2 shows that the POD projection error decreases rapidly for both outputs. In the displacement assessment, POD-GPR and PODI-RBF most closely follow the decreasing projection error before reaching regression-controlled plateaus, whereas PODI-Linear and the neural coefficient maps show stronger dimension-dependent saturation. At comparable retained dimensions, the thermal-driver projection error is several orders of magnitude smaller than the displacement projection error; equivalently, fewer POD modes are needed to represent the thermal-driver snapshots at the same accuracy. PODI-RBF and POD-GPR also give the lowest predictive thermal-driver validation errors over the tested dimensions. The non-monotone POD-NN, POD-DL-ROM, and DL-ROM curves show that increasing the reduced or latent dimension alone does not guarantee better validation accuracy.

These validation results support the common working choices $N_w = 6$ for displacement and $N_\theta = 4$ for the thermal driver in the three-case comparison. They are common comparison dimensions rather than method-specific validation minimizers. At these dimensions, Figure 8.1 confirms that $\mathcal{E}^{\chi}_{N_\chi} \geq 99.999\%$ for every campaign. The minimum fractions of the computed 30-mode spectrum are 99.999707% for displacement and 99.99999987% for the thermal driver.

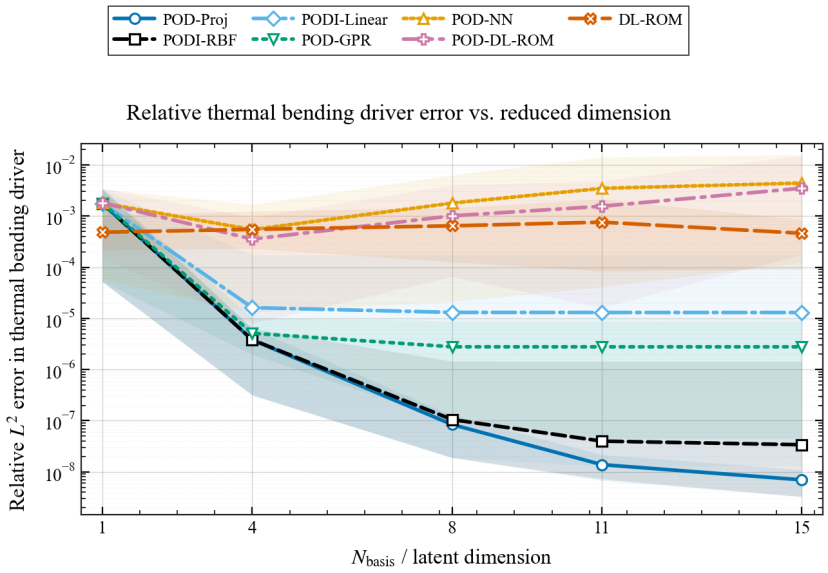
After the dimensions have been fixed, the held-out POD-projection errors provide a representation diagnostic. Their mean values are approximately $7.95\times10^{-3}\%$, $1.08\times10^{-1}\%$, and $2.13\times10^{-1}\%$ for the displacement fields in Cases 2, 3, and 5, respectively. The corresponding thermal-driver values are approximately $6.76\times10^{-4}\%$, $1.30\times10^{-3}\%$, and $3.47\times10^{-3}\%$. The larger predictive errors observed in the more demanding parameter domains therefore arise primarily from approximating the parameter-to-output map rather than from POD truncation.

**Hybrid POD–Galerkin verification.** A representative hybrid calculation is performed using the selected displacement dimension $N = 6$. At the selected validation parameter, the direct POD projection gives a relative displacement error of 0.0773%, whereas the thermomechanical POD–Galerkin iteration gives 0.1119%. The projected calculation converges in 11 partitioned iterations, with final displacement and thermal-driver update indicators $2.934 \times 10^{-6}$ and $7.400 \times 10^{-8}$, respectively. The POD–Galerkin error is slightly larger than the corresponding projection benchmark, as expected for a reduced equilibrium solve.

The thermal problem still contains 36,465 degrees of freedom, and the plate thermal-driver space contains 5,194 degrees of freedom. Moreover,




**(a) Displacement-only $w$**




**(b) Thermal-driver-only $\theta$**

**Figure 8.2:** Case 2 validation error versus reduced or latent dimension for (a) displacement and (b) thermal driver. Curves show the mean relative finite element $L^2(\Omega)$ error, while the shaded bands show the samplewise minimum–maximum range. POD-Proj is the non-predictive projection baseline.

the mechanical residual and Jacobian are assembled in the full-order space before projection. The experiment therefore verifies the consistency of the projected mechanical correction and its interaction with the partitioned iteration; it does not represent a completely reduced intrusive thermomechanical online solver.

**Held-out predictive accuracy across Cases 2, 3, and 5.** After the reduced dimensions and model checkpoints have been fixed, the models are evaluated on the case-specific test subsets: 45 samples in Case 2, 38 samples in Case 3, and 50 samples in Case 5. Figure 8.3 consolidates the final field-wise errors in a $2 \times 3$ arrangement. Columns correspond to the three physical cases, and rows correspond to the displacement-only and thermal-driver-only outputs.

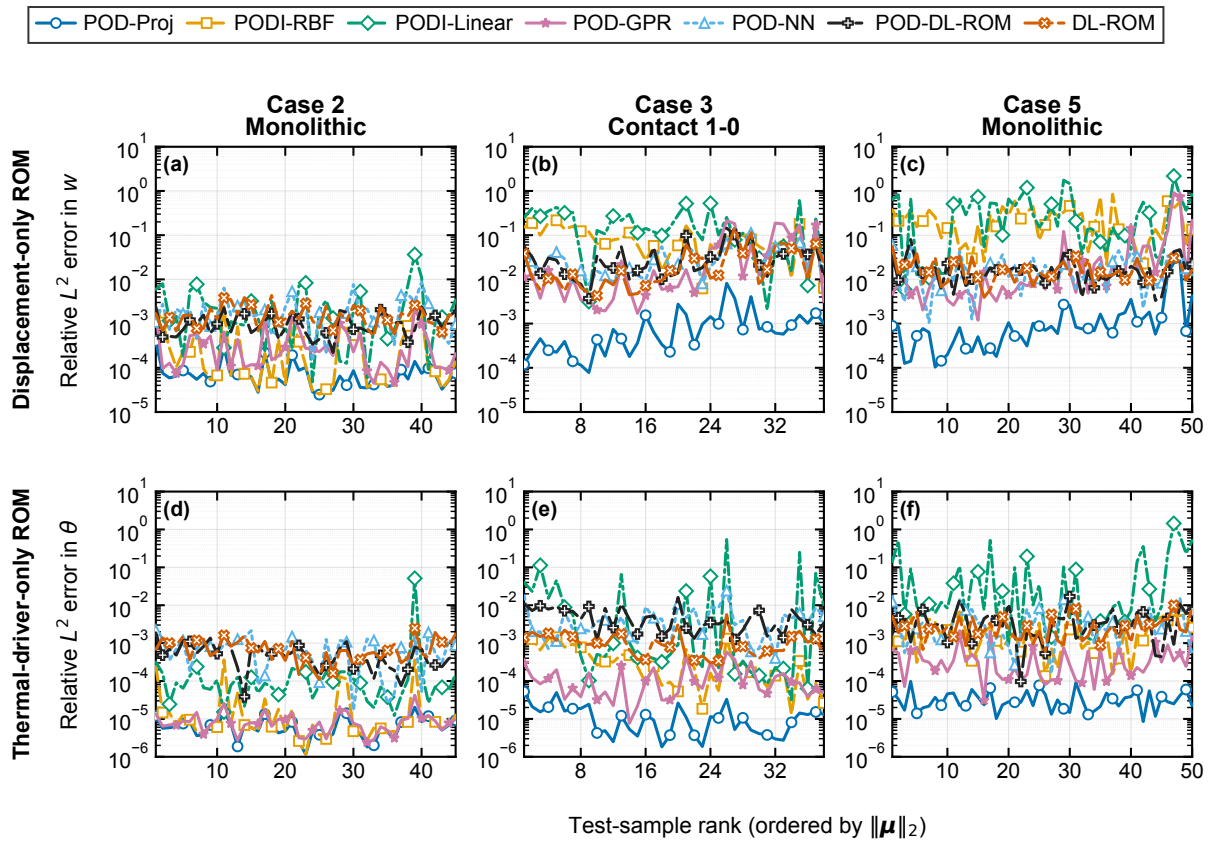



**Figure 8.3:** Test-set relative $L^2(\Omega)$ errors for Cases 2, 3, and 5 (columns) and the displacement-only and thermal-driver-only outputs (rows). The subsets contain 45, 38, and 50 samples, ordered by $\|\mu\|_2$. POD-Proj is the non-predictive projection baseline; the fixed dimensions are $N_w = 6$ and $N_\theta = 4$.

Table 8.4 extracts the lowest mean predictive error for each one-field output and physical case. It provides a compact quantitative summary of Figure 8.3; the complete samplewise curves remain essential for identifying outliers and variations over the parameter domain.

The Case 2 parameter-to-field maps are sufficiently smooth that POD-GPR and PODI-RBF remain close to the projection baseline for both outputs. POD-GPR gives the lowest mean predictive error for the displacement and thermal-driver assessments. The direct and POD-enhanced neural models remain accurate, but their errors do not decrease as rapidly as the projection baseline in this low-dimensional parameter domain.

The Case 3 contact configuration and the five-parameter Case 5 configuration produce more difficult displacement maps. In these cases, DL-ROM gives the lowest mean test-set displacement error. POD-NN and POD-DL-ROM are also competitive, particularly in the Case 5 displacement study. POD-GPR often retains a low median displacement error, but isolated larger errors increase its mean and maximum statistics in the higher-dimensional cases.

The thermal-driver response remains substantially easier to predict than the displacement response. POD-GPR gives the lowest mean thermal-driver error in all three cases, and PODI-RBF is generally the next most accurate coefficient map. The isolated large errors produced by PODI-Linear, especially in Cases 3 and 5, indicate sensitivity to sparse local support in the higher-dimensional parameter domains. These

**Table 8.4:** Lowest mean test-set relative $L^2(\Omega)$ error (%) among predictive one-field ROMs; POD-Proj is excluded.

| | Case 2 | | Case 3 | | Case 5 | |
|---|---|---|---|---|---|---|
| output | method | mean error | method | mean error | method | mean error |
| displacement-only $w$ | POD-GPR | $4.10 \times 10^{-2}$ | DL-ROM | $2.748 \times 10^{0}$ | DL-ROM | $1.733 \times 10^{0}$ |
| thermal-driver-only $\theta$ | POD-GPR | $9.25 \times 10^{-4}$ | POD-GPR | $1.33 \times 10^{-2}$ | POD-GPR | $4.22 \times 10^{-2}$ |

results also demonstrate why projection error alone is not sufficient for ROM assessment: in the more demanding displacement problems, the dominant error arises from learning the parameter-to-coordinate map rather than from the six-mode POD truncation.

**Online computational performance.** Figure 8.4 reports the samplewise timing ratios for the same held-out test subsets used in the accuracy assessment. The columns correspond to Cases 2, 3, and 5, with 45, 38, and 50 test samples, respectively, and the rows correspond to the displacement and thermal-driver outputs. All methods use the selected dimensions $N_w = 6$ and $N_\theta = 4$. For a predictive method, the plotted quantity is $S^{(j)}_{\chi,m}$ from Equation (8.21). POD-Proj instead reports $R^{(j)}_{\chi,\mathrm{proj}}$, which requires an available full-order field and is therefore only a projection-cost diagnostic.

Across the six panels, the predictive speedups span approximately $10^3$ to $10^6$. The displacement-only models principally occupy the range $10^3$–$10^5$, with the direct full-state DL-ROM generally near the lower end because of its more expensive decoding and field reconstruction. The thermal-driver-only models lie predominantly in the range $10^5$–$10^6$, and the coefficient-based methods generally occupy the upper part of that range. The samplewise variation reflects both variations in the coupled FOM wall time and method-specific differences in predictive evaluation and reconstruction cost. These quantities are matched empirical online ratios under the reported timing definition rather than hardware-independent complexity estimates.

**Spatial reconstruction and pointwise error diagnostics.** The samplewise relative errors in Figure 8.3 provide the complete cross-method comparison over the test subsets. They are complemented first by the displacement projection fields in Figure 8.5, which show the retained six-mode displacement representation and its pointwise truncation error at one preselected test sample per case.

Figure 8.5 shows that the six-mode POD projection reproduces the dominant displacement pattern in all three cases. The remaining truncation error is small relative to the field magnitude and is concentrated in regions where the displacement varies most strongly.

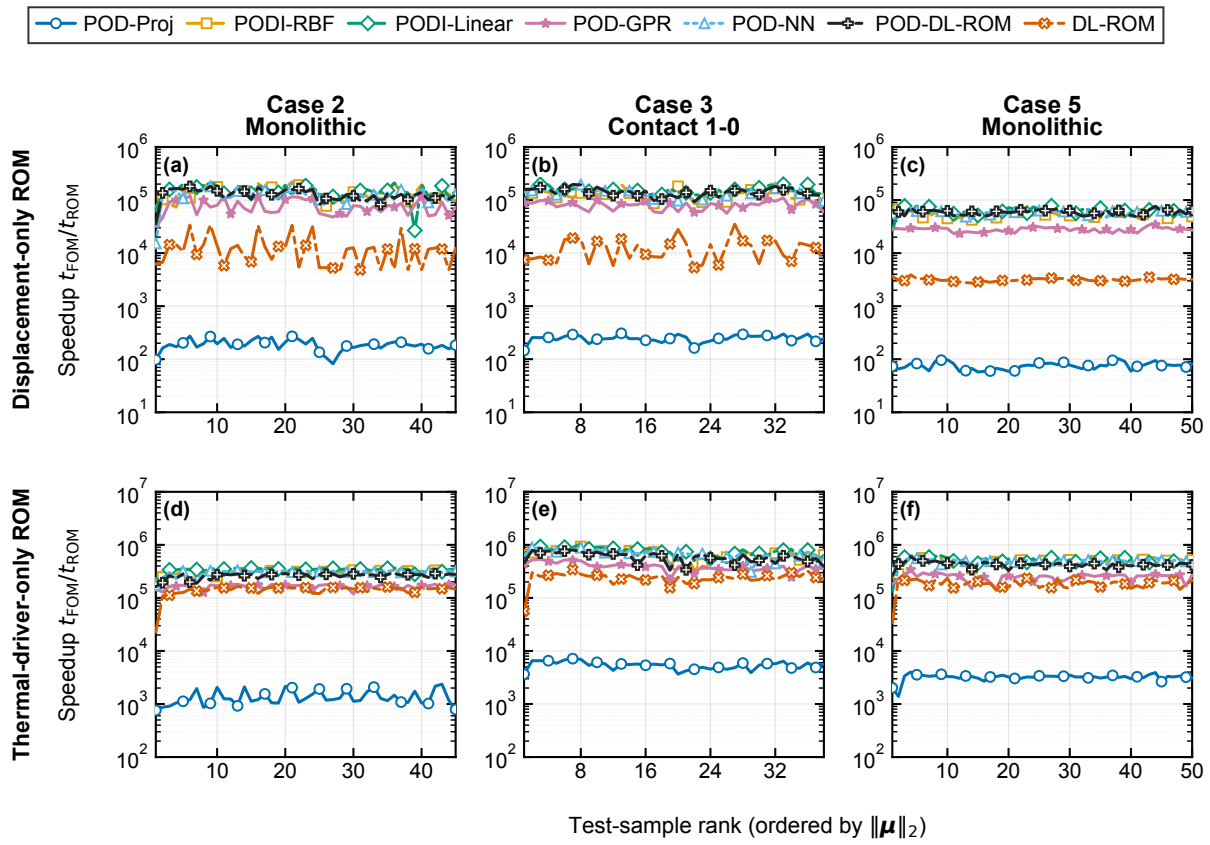



**Figure 8.4:** Samplewise timing ratios for Cases 2, 3, and 5. Columns correspond to the three cases and rows to the displacement and thermal-driver outputs. Predictive curves report $S^{(j)}_{\chi,m}$; POD-Proj reports the non-predictive projection-cost ratio $R^{(j)}_{\chi,\text{proj}}$.

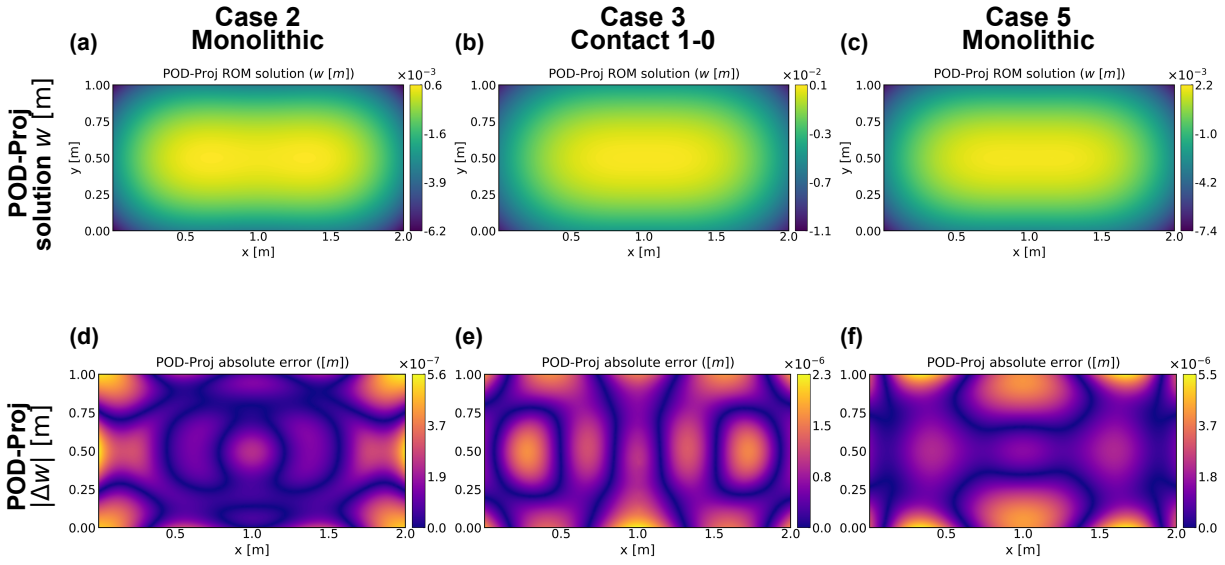



**Figure 8.5:** POD-Proj displacement (top) and pointwise absolute projection error (bottom) for preselected test samples from Cases 2, 3, and 5. The same case-specific displacement sample is used in both rows; color scales are panel-specific. Aggregate projection and predictive accuracy are reported in Figure 8.3 and Table 8.4.

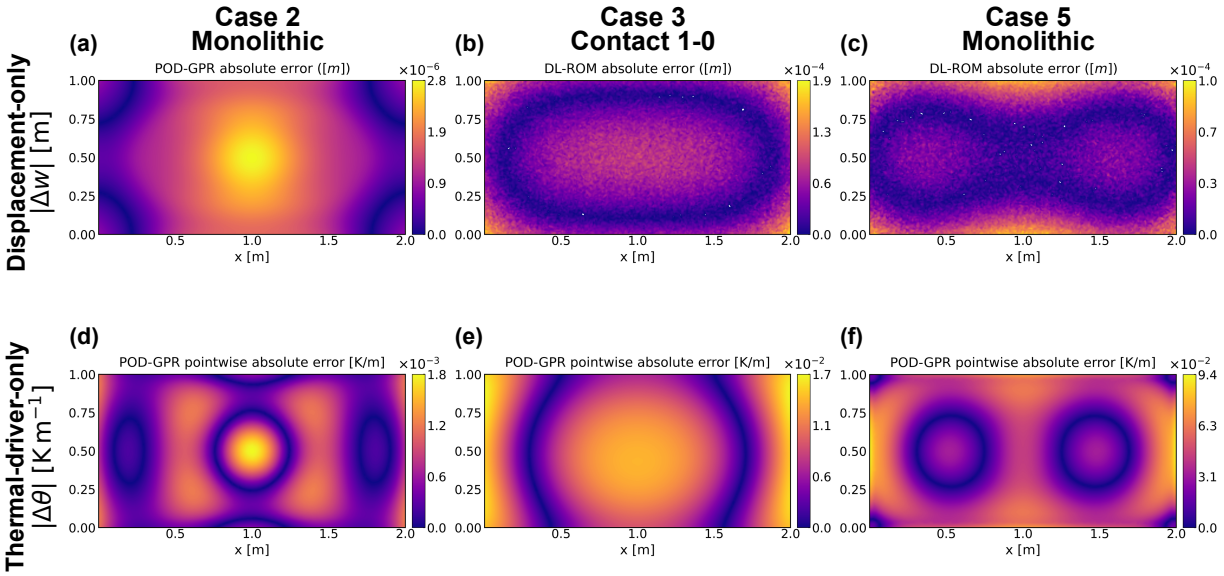



**Figure 8.6:** Pointwise absolute-error fields for preselected test samples from Cases 2, 3, and 5 (columns) and the displacement-only and thermal-driver-only outputs (rows). For each panel, the displayed method is the predictive ROM with the lowest mean error in Table 8.4. Color scales are independent; displacement and thermal-driver errors are reported in meters and K m$^{-1}$, respectively.

Figure 8.6 then reports the selected predictive absolute-error fields for the two retained outputs. For each case and output, the displayed method is the method attaining the lowest mean relative $L^2(\Omega)$ error on the corresponding test subset in Table 8.4. Accordingly, POD-GPR is shown for both Case 2 outputs and for the thermal-driver outputs in Cases 3 and 5, whereas DL-ROM is shown for the displacement outputs in Cases 3 and 5. The method selection follows the aggregate test-set statistics rather than the particular field sample displayed.

The spatial maps support the trends observed in the aggregate test-set assessment. The Case 2 errors are comparatively small and spatially smooth for both retained outputs. The Case 3 contact configuration and the five-parameter Case 5 configuration produce more heterogeneous displacement errors, consistent with their more demanding parameter-to-field maps and the higher relative errors reported in Table 8.4. The thermal-driver error maps remain spatially structured, and their lower relative $L^2(\Omega)$ errors are consistent with the aggregate statistics.

## 8.8 Interpretation, limitations, and chapter conclusions

The POD spectra show that both retained thermomechanical outputs admit low-dimensional linear representations. The final displacement comparisons use six POD modes, while the thermal-driver comparisons use four. These spaces represent displacement vectors containing approximately $2 \times 10^4$ components and thermal-driver vectors containing approximately $5 \times 10^3$ components. At the selected dimensions, the bases account for at least 99.999% of the sum of the first 30 computed POD eigenvalues for the corresponding field in every case, while preserving field-specific metrics and physical units.

The hybrid POD–Galerkin verification shows that the mechanical correction can be projected without disrupting the partitioned convergence

behavior. The retained three-dimensional thermal solve and full-order residual and Jacobian assembly nevertheless make this a hybrid reduction. A fully reduced intrusive thermomechanical implementation would additionally require reduced three-dimensional thermal states, reduced thermal-to-plate transfer operators, and efficient treatment of the nonlinear and parameter-dependent operators.

The non-intrusive models eliminate the partitioned online iteration for the selected output. Under the matched single-query timing definition in Equation (8.21), the predictive models achieve samplewise speedups of approximately $10^3$–$10^6$ across the three cases, with the thermal-driver-only models generally occupying the upper part of that range. POD-GPR gives the lowest mean errors for all thermal-driver outputs and for both Case 2 assessments. In the higher-dimensional Case 3 and Case 5 displacement problems, DL-ROM gives the lowest mean test-set errors, with POD-NN and POD-DL-ROM remaining competitive. PODI-RBF remains effective in the smooth low-dimensional setting, whereas its relative performance varies more strongly in the higher-dimensional cases. PODI-Linear exhibits the largest isolated errors in those cases, indicating the limitations of simplicial interpolation when local parameter-space support is sparse.

The comparison with the POD projection baseline shows that the remaining displacement error in Cases 3 and 5 is governed primarily by the learned parameter-to-coordinate or parameter-to-state map rather than by six-mode POD truncation. The field-wise POD spectra and dimension studies establish the low-dimensional representation properties of the two outputs, while the test-set comparisons and spatial error maps quantify predictive accuracy and reconstruction quality across the three cases.

The two one-field surrogates are valid within the sampled steady thermomechanical domains. Their predictive accuracy depends on parameter-space coverage, and the POD–GPR posterior standard deviations are coordinate-level model-uncertainty indicators.

Chapter 9 next turns from steady thermomechanical bending to time-dependent transverse dynamics. The transient study retains the common orthotropic plate and panel-interface framework while introducing inertia, time-dependent forcing, and trajectory-level reduced outputs under a prescribed foundation regime.

# Parameterized Transient Dynamics of Orthotropic Kirchhoff–Love Plates

# 9


9.1 Role and scope of the transient model . . . . . 102
9.2 Geometry, panel partition, and notation . . . . . 102
9.3 Kirchhoff–Love kinematics and orthotropic bending 104
9.4 Kinetic, elastic, foundation, interface, and loading contributions . . . . . 105
9.5 Hamilton's principle and weak evolution problem . 107
9.6 Strong form, boundary conditions, and initial data 108
9.7 Parameterized transient solution map . . . . . 109
9.8 Summary and handoff to the full-order approximation . . . . . 110


## 9.1 Role and scope of the transient model

The preceding chapters develop the static mechanical and thermomechanical responses of parameterized orthotropic Kirchhoff–Love plates. This chapter introduces the transient-dynamics study: the free-vibration, forced-response, and transient transverse behavior of monolithic and panel-partitioned plates on elastic Winkler foundations. The formulation preserves the thesis-wide upward-positive sign convention, aligned orthotropic bending law, and weak transmission of displacement and normal rotation across physical panel interfaces. Its additional ingredients are transverse inertia, initial displacement and velocity, time-dependent loading, and a parameter–time solution map.

The foundation treatment is more restricted than the unilateral static model of Chapter 3. One effective Winkler coefficient is selected before a trajectory is advanced and remains fixed throughout the time interval. The evolution problem is therefore linear for each prescribed support regime. It supports free-vibration, harmonic-response, and transient analyses, while time-dependent detachment, recontact, impact, and compression–uplift switching would require a different state-switching evolution model. This distinction is retained in the full-order and reduced-order formulations and is essential for interpreting the reported results.

The model retains transverse translational inertia only. Rotary inertia, transverse-shear deformation, in-plane kinematics and inertia, geometric nonlinearity, thermomechanical coupling, and state-dependent damping remain outside the governing formulation [94]. Optional Rayleigh damping is introduced later at the semidiscrete algebraic level, whereas the transient campaigns used for model reduction are undamped. These are assumptions of the parent transient model and are inherited unchanged by its reduced-order approximations.

The chapter first fixes the geometric, kinematic, energetic, and sign conventions. It then derives the weak initial–boundary-value problem, records the corresponding strong form and boundary conditions, and defines the parameterized transient solution map used in the many-query setting. Chapter 10 develops and verifies the full-order approximation, while Chapter 11 constructs and assesses the non-intrusive reduced models.

## 9.2 Geometry, panel partition, and notation

Let $T > 0$ be the final time and let the plate midsurface occupy a bounded Lipschitz domain $\Omega \subset \mathbb{R}^2$. Although the variational formulation is not restricted to rectangular geometries, the numerical campaigns use

$$\Omega = (0, L) \times (0, B), \qquad L = 2\,\mathrm{m}, \qquad B = 1\,\mathrm{m}. \tag{9.1}$$

The plate has constant thickness $t > 0$, and its three-dimensional reference body is $\Omega \times (-t/2, t/2)$. Let $(\mathbf{e}_x, \mathbf{e}_y, \mathbf{e}_z)$ be an orthonormal basis of $\mathbb{R}^3$, with

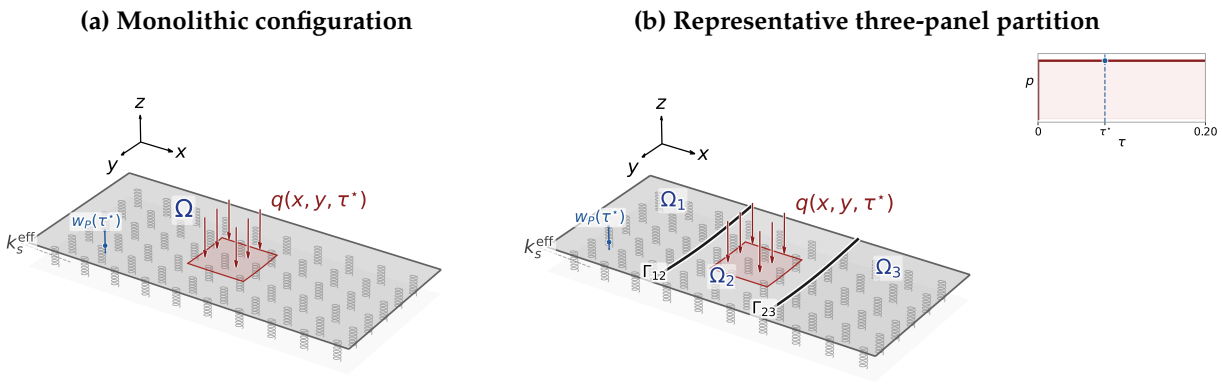



**Figure 9.1:** Geometric setting for the transient plate model at a representative time $\tau^\star \in [0, T]$. In both configurations, the plate rests on the prescribed effective Winkler coefficient $k_s^{\text{eff}}$, is subjected to the localized transverse load $q(x, y, \tau^\star)$, and is observed through the probe displacement $w_P(\tau^\star)$. Panel (a) shows the monolithic midsurface $\Omega$. Panel (b) shows the schematic partition $\overline{\Omega} = \overline{\Omega}_1 \cup \overline{\Omega}_2 \cup \overline{\Omega}_3$ with interfaces $\Gamma_{12}$ and $\Gamma_{23}$; the inset identifies the step-type loading profile. The continuous model permits any partition satisfying Equation (9.2); the reported panel ROM campaign uses the two-panel configuration defined in Section 11.3.

associated Cartesian coordinates $(x, y, z)$ and $z = 0$ on the midsurface. The material orthotropy directions are aligned with the $x$- and $y$-axes.

A modular assembly is represented by a non-overlapping partition into $M$ open Lipschitz subdomains,

$$\overline{\Omega} = \bigcup_{i=1}^{M} \overline{\Omega_i}, \qquad \Omega_i \cap \Omega_j = \varnothing \quad \text{for } i \neq j. \tag{9.2}$$

For two adjacent panels, define the physical interface and the complete interface skeleton by

$$\Gamma_{ij} := \partial\Omega_i \cap \partial\Omega_j, \qquad \Gamma := \bigcup_{1 \le i < j \le M} \Gamma_{ij}. \tag{9.3}$$

On each $\Gamma_{ij}$, choose once and for all the reference unit normal $\mathbf{n} := \mathbf{n}_i = -\mathbf{n}_j$, where $\mathbf{n}_i$ and $\mathbf{n}_j$ are the outward normals of $\Omega_i$ and $\Omega_j$, respectively. For any panel-wise scalar field $v = \{v_i\}_{i=1}^{M}$, set

$$\partial_n v_i := \mathbf{n} \cdot \nabla v_i, \qquad [\![v]\!]_{ij} := v_i - v_j, \qquad [\![\partial_n v]\!]_{ij} := \partial_n v_i - \partial_n v_j \quad \text{on } \Gamma_{ij}. \tag{9.4}$$

Thus, $[\![v]\!]_{ij}$ is the displacement jump and $[\![\partial_n v]\!]_{ij}$ is the jump of the normal rotation under the adopted Kirchhoff–Love convention.

The exterior boundary is decomposed into disjoint relatively open portions $\Gamma_S, \Gamma_F \subset \partial\Omega$ such that

$$\partial\Omega = \overline{\Gamma_S} \cup \overline{\Gamma_F}, \qquad \Gamma_S \cap \Gamma_F = \varnothing, \tag{9.5}$$

where $\Gamma_S$ and $\Gamma_F$ denote simply supported and free portions, respectively. On a smooth exterior edge, $\mathbf{n} = (n_x, n_y)$ denotes the outward unit normal, $\mathbf{t} = (-n_y, n_x)$ the corresponding unit tangent, and $\partial_s := \mathbf{t} \cdot \nabla$ the tangential derivative.

The unknown is the panel-wise transverse displacement

$$w = \{w_i\}_{i=1}^{M}, \qquad w_i : \Omega_i \times [0, T] \to \mathbb{R}. \tag{9.6}$$

Time is denoted by $\tau$, with $\dot{w} := \partial_\tau w$ and $\ddot{w} := \partial_{\tau\tau} w$ for the transverse velocity and acceleration. Displacement and transverse load are positive in the $\mathbf{e}_z$-direction; consequently, a downward load has negative sign. For

a selected observation point $P \in \Omega \setminus \Gamma$, the probe history is denoted by

$$w_P(\tau) := w(P, \tau). \tag{9.7}$$

## 9.3 Kirchhoff–Love kinematics and orthotropic bending

Under the small-deflection Kirchhoff–Love assumptions, the displacement of a material point $(x, y, z)$ in panel $\Omega_i$ is

$$\begin{aligned} &u_x(x, y, z, \tau) = -z\, \partial_x w_i(x, y, \tau), \quad u_y(x, y, z, \tau) = -z\, \partial_y w_i(x, y, \tau), \\ &u_z(x, y, z, \tau) = w_i(x, y, \tau). \end{aligned} \tag{9.8}$$

Hence, fibers initially normal to the midsurface remain straight and normal after deformation, and the transverse-shear strains vanish. The implemented dynamic model retains only the translational inertia associated with $u_z$. The rotary-inertia contribution generated by $\dot{u}_x$ and $\dot{u}_y$ is not included. In the notation of classical plate dynamics, the retained inertia is therefore the mass per unit area

$$I_0 = \rho t, \tag{9.9}$$

whereas the rotary-inertia coefficient

$$I_2 = \frac{\rho t^3}{12} \tag{9.10}$$

is omitted [12–14].

For each panel, introduce the engineering-curvature vector

$$\boldsymbol{\kappa}(w_i) := \begin{bmatrix} \partial_{xx} w_i \\ \partial_{yy} w_i \\ 2\partial_{xy} w_i \end{bmatrix} \tag{9.11}$$

and the aligned-orthotropic rigidity matrix

$$\mathbf{D}(\boldsymbol{\mu}) := \begin{bmatrix} D_x & D_{xy} & 0 \\ D_{xy} & D_y & 0 \\ 0 & 0 & D_s \end{bmatrix}. \tag{9.12}$$

The flexural and torsional rigidities have units of force times length and are assumed to satisfy

$$D_x > 0, \qquad D_y > 0, \qquad D_s > 0, \qquad D_x D_y - D_{xy}^2 > 0, \tag{9.13}$$

which ensures positive definiteness of the bending energy.

For consistency with the finite element formulation, define the curvature-conjugate moment operators

$$\begin{aligned} M_x(w_i) &:= D_x \partial_{xx} w_i + D_{xy} \partial_{yy} w_i, \\ M_y(w_i) &:= D_{xy} \partial_{xx} w_i + D_y \partial_{yy} w_i, \\ M_{xy}(w_i) &:= 2 D_s \partial_{xy} w_i. \end{aligned} \tag{9.14}$$

With the upward-positive displacement convention in Equation (9.8), the

corresponding physical bending-moment tensor is

$$\mathbf{M}^{\mathrm{phys}}(w_i) = -\begin{bmatrix} M_x(w_i) & M_{xy}(w_i) \\ M_{xy}(w_i) & M_y(w_i) \end{bmatrix}. \tag{9.15}$$

This explicit distinction reconciles the positive curvature-conjugate operator used in the energy and $C^0$ interior-penalty terms with the physical sign of the moment resultants. Since the exterior moment conditions are homogeneous, they are invariant under this global sign change.

The factor 2 in $M_{xy}$ is essential. With engineering curvature $2w_{xy}$, the twisting contribution to the energy is $2D_s w_{xy}^2$, and its associated bilinear contribution is $4D_s w_{xy} v_{xy}$.

## 9.4 Kinetic, elastic, foundation, interface, and loading contributions

Let $\rho > 0$ denote the mass density. Consistently with the transverse-only inertia retained in the implementation, the kinetic energy is

$$\mathcal{T}(\dot{w}) := \frac{1}{2}\sum_{i=1}^{M}\int_{\Omega_i} \rho t\, \dot{w}_i^2 \,\mathrm{d}\Omega. \tag{9.16}$$

The orthotropic bending energy is

$$\begin{aligned} \mathcal{U}_b(w;\boldsymbol{\mu}) &:= \frac{1}{2}\sum_{i=1}^{M}\int_{\Omega_i} \boldsymbol{\kappa}(w_i)^{\mathsf{T}}\mathbf{D}(\boldsymbol{\mu})\boldsymbol{\kappa}(w_i)\,\mathrm{d}\Omega \\ &= \frac{1}{2}\sum_{i=1}^{M}\int_{\Omega_i} \Big[ D_x(\partial_{xx}w_i)^2 + D_y(\partial_{yy}w_i)^2 \\ &\qquad + 2D_{xy}(\partial_{xx}w_i)(\partial_{yy}w_i) \\ &\qquad + 4D_s(\partial_{xy}w_i)^2 \Big]\,\mathrm{d}\Omega. \end{aligned} \tag{9.17}$$

The static mechanical and thermomechanical models distinguish compression and uplift through a unilateral foundation law. In the present transient implementation, however, one effective regime is selected before a trajectory is advanced and remains fixed over the complete time interval. Let $k_s \geq 0$ denote the Winkler modulus, with units of force per unit volume [15]. The prescribed effective Winkler coefficient is

$$k_s^{\mathrm{eff}}(\boldsymbol{\mu}) = \begin{cases} 0, & \text{no foundation,} \\ k_s, & \text{prescribed compression regime,} \\ \varepsilon_f k_s, & \text{prescribed regularized-uplift regime.} \end{cases} \tag{9.18}$$

Here $0 < \varepsilon_f \ll 1$. Its potential energy is

$$\mathcal{U}_f(w;\boldsymbol{\mu}) := \frac{1}{2}\sum_{i=1}^{M}\int_{\Omega_i} k_s^{\mathrm{eff}}(\boldsymbol{\mu}) w_i^2 \,\mathrm{d}\Omega. \tag{9.19}$$

Equation (9.18) must not be interpreted as a state-dependent contact law: $k_s^{\mathrm{eff}}$ does not switch when $w_i$ changes sign. All three sampled snapshot

campaigns use the prescribed compression regime,

$$k_s^{\mathrm{eff}} = k_s. \tag{9.20}$$

Mechanical transmission across physical panel seams is represented by the penalty energy

$$\mathcal{U}_\Gamma(w;\boldsymbol{\mu}) := \frac{1}{2} \sum_{1\le i<j\le M} \int_{\Gamma_{ij}} \left[ \eta_w \, [\![w]\!]_{ij}^2 + \eta_\psi \, [\![\partial_n w]\!]_{ij}^2 \right] \mathrm{d}\Gamma, \tag{9.21}$$

where $\eta_w > 0$ and $\eta_\psi > 0$ are the displacement- and normal-rotation-penalty weights. Since $[\![w]\!]_{ij}$ has units of length and $[\![\partial_n w]\!]_{ij}$ is dimensionless, $\eta_w$ and $\eta_\psi$ have units of force per unit area and force, respectively. For a monolithic plate, $M = 1$ and $\mathcal{U}_\Gamma = 0$. At finite weights, Equation (9.21) defines the penalty-coupled panel model; formally, increasing the weights suppresses the kinematic jumps $[\![w]\!]_{ij}$ and $[\![\partial_n w]\!]_{ij}$. These numerical weights reproduce the reported penalty-coupled model; they are not interpreted as experimentally calibrated connector stiffnesses.

The transverse load is written as the product of a spatial distribution and a dimensionless temporal factor,

$$q(x, y, \tau; \boldsymbol{\mu}) = \widehat{q}(x, y; \boldsymbol{\mu})\, s(\tau; \boldsymbol{\mu}). \tag{9.22}$$

The corresponding external-work functional is

$$\ell_{\boldsymbol{\mu}}(\tau; v) := \sum_{i=1}^{M} \int_{\Omega_i} q(x, y, \tau; \boldsymbol{\mu}) v_i \, \mathrm{d}\Omega. \tag{9.23}$$

For the active load-amplitude campaign, the spatial factor is a rectangular patch centered at

$$(x_c, y_c) = (1\,\mathrm{m}, 0.5\,\mathrm{m}) \tag{9.24}$$

with side lengths

$$a_x = a_y = 0.351\,\mathrm{m}. \tag{9.25}$$

More precisely,

$$\begin{aligned} P_q &:= \left[ x_c - \frac{a_x}{2}, x_c + \frac{a_x}{2} \right] \times \left[ y_c - \frac{a_y}{2}, y_c + \frac{a_y}{2} \right], \\ \widehat{q}(x, y; A_q) &:= -A_q \, \chi_{P_q}(x, y), \qquad A_q > 0, \end{aligned} \tag{9.26}$$

where $A_q$ is the absolute load amplitude in force per unit area and the minus sign produces downward loading. The snapshot generators use $a_x = a_y = 0.351\,\mathrm{m}$; the separate localized-load verification in Section 10.6 uses the independently specified value $0.35\,\mathrm{m}$ and the positive pressure $1500\,\mathrm{N\,m^{-2}}$ under the upward-positive sign convention. These values are therefore reported without silently identifying the two configurations.

The temporal profile used in the reported campaign is the step

$$s(\tau) = H(\tau) = 1 \qquad \text{for } 0 \le \tau \le T, \tag{9.27}$$

with the implemented convention $H(0) = 1$. Thus, zero initial displacement and velocity are compatible with a generally nonzero initial acceleration.

## 9.5 Hamilton's principle and weak evolution problem

Define the broken mass and energy spaces

$$\begin{aligned} H &:= \prod_{i=1}^{M} L^2(\Omega_i), \\ V &:= \left\{ v = \{v_i\}_{i=1}^{M} : v_i \in H^2(\Omega_i),\ v_i = 0 \text{ on } \Gamma_S \cap \partial\Omega_i \right\}. \end{aligned} \tag{9.28}$$

The total elastic energy is

$$\mathcal{U}(w;\boldsymbol{\mu}) := \mathcal{U}_b(w;\boldsymbol{\mu}) + \mathcal{U}_f(w;\boldsymbol{\mu}) + \mathcal{U}_\Gamma(w;\boldsymbol{\mu}). \tag{9.29}$$

For admissible variations $v$ that vanish at the temporal endpoints, Hamilton's principle reads [95]

$$\delta \int_0^T \left[ \mathcal{T}(\dot{w}) - \mathcal{U}(w;\boldsymbol{\mu}) + \ell_{\boldsymbol{\mu}}(\tau; w) \right] \mathrm{d}\tau = 0. \tag{9.30}$$

After integration by parts in time, the parameterized weak problem is: for $\boldsymbol{\mu} \in \mathcal{D}$, find $w(\tau;\boldsymbol{\mu}) \in V$ such that

$$m(\ddot{w}, v) + a_{\boldsymbol{\mu}}^{\mathrm{eff}}(w, v) + c_{\Gamma,\boldsymbol{\mu}}(w, v) = \ell_{\boldsymbol{\mu}}(\tau; v) \qquad \forall v \in V. \tag{9.31}$$

The identity holds for almost every $\tau \in (0, T)$. The mass form is

$$m(u, v) := \sum_{i=1}^{M} \int_{\Omega_i} \rho t\, u_i v_i \,\mathrm{d}\Omega. \tag{9.32}$$

The plate-and-foundation form is

$$\begin{aligned} a_{\boldsymbol{\mu}}^{\mathrm{eff}}(w, v) := \sum_{i=1}^{M} \int_{\Omega_i} \Big[ & D_x(\partial_{xx} w_i)(\partial_{xx} v_i) + D_y(\partial_{yy} w_i)(\partial_{yy} v_i) \\ & + D_{xy}\big((\partial_{xx} w_i)(\partial_{yy} v_i) + (\partial_{yy} w_i)(\partial_{xx} v_i)\big) \\ & + 4D_s(\partial_{xy} w_i)(\partial_{xy} v_i) + k_s^{\mathrm{eff}} w_i v_i \Big] \,\mathrm{d}\Omega. \end{aligned} \tag{9.33}$$

The panel-interface form is

$$\begin{aligned} c_{\Gamma,\boldsymbol{\mu}}(w, v) := \sum_{1 \le i < j \le M} \int_{\Gamma_{ij}} \Big[ & \eta_w \, [\![ w ]\!]_{ij} \, [\![ v ]\!]_{ij} \\ & + \eta_\psi \, [\![ \partial_n w ]\!]_{ij} \, [\![ \partial_n v ]\!]_{ij} \Big] \,\mathrm{d}\Gamma. \end{aligned} \tag{9.34}$$

Whenever the combined stiffness form is coercive on $V$, or, for an unsupported plate with entirely free edges, after the finite-dimensional rigid-motion kernel has been treated separately, the standard linear second-order evolution theory applies [96]. For $w_0 \in V$, $v_0 \in H$, and $\ell_{\boldsymbol{\mu}} \in L^2(0, T; V')$, the natural energy regularity is

$$w \in C([0, T]; V), \qquad \dot{w} \in C([0, T]; H), \qquad \ddot{w} \in L^2(0, T; V'). \tag{9.35}$$

This formulation gives a direct meaning to the initial conditions stated below. A separate existence-and-uniqueness theorem is not stated here for the penalty-coupled evolution problem.

The continuous problem (9.31) is undamped. This is deliberate: the implementation introduces optional Rayleigh damping only after spatial discretization through the standard proportional-damping form [52, 95]:

$$\mathbf{C} = \alpha_M \mathbf{M} + \alpha_K \mathbf{K}. \tag{9.36}$$

Accordingly, no continuum viscous constitutive operator is assigned to this algebraic damping choice. The active snapshot campaigns use

$$\alpha_M = 0\,\mathrm{s}^{-1}, \qquad \alpha_K = 0\,\mathrm{s}. \tag{9.37}$$

The units follow from $\mathbf{C} = \alpha_M \mathbf{M} + \alpha_K \mathbf{K}$.

## 9.6 Strong form, boundary conditions, and initial data

For constant rigidities and a sufficiently smooth solution in the interior of a panel, Equation (9.31) gives

$$\begin{aligned}\rho t\, \partial_{\tau\tau} w_i + D_x \partial_{xxxx} w_i + 2(D_{xy} + 2D_s)\partial_{xxyy} w_i \\ + D_y \partial_{yyyy} w_i + k_s^{\mathrm{eff}} w_i = q \qquad \text{in } \Omega_i \times (0, T).\end{aligned} \tag{9.38}$$

The mixed-derivative coefficient follows from the two $D_{xy}$ contributions and the twisting operator $M_{xy} = 2D_s w_{xy}$. Equation (9.38) is the aligned-orthotropic classical plate equation with translational inertia $\rho t$ and without rotary inertia.

To state the exterior natural conditions in the same variational sign convention as Equation (9.14), define

$$M_{nn}(w) := M_x(w) n_x^2 + 2M_{xy}(w) n_x n_y + M_y(w) n_y^2, \tag{9.39}$$

$$M_{nt}(w) := (M_y(w) - M_x(w)) n_x n_y + M_{xy}(w)(n_x^2 - n_y^2), \tag{9.40}$$

$$\begin{aligned} Q_x(w) &:= \partial_x M_x(w) + \partial_y M_{xy}(w), \\ Q_y(w) &:= \partial_x M_{xy}(w) + \partial_y M_y(w), \\ Q_n(w) &:= Q_x(w) n_x + Q_y(w) n_y, \end{aligned} \tag{9.41}$$

$$V_n(w) := Q_n(w) + \partial_s M_{nt}(w). \tag{9.42}$$

The corresponding physical resultants are the negatives of these quantities, consistent with Equation (9.15); the homogeneous boundary conditions are unchanged. On a simply supported edge,

$$w = 0, \qquad M_{nn}(w) = 0 \quad \text{on } \Gamma_S. \tag{9.43}$$

On a free edge,

$$M_{nn}(w) = 0, \qquad V_n(w) = 0 \quad \text{on } \Gamma_F. \tag{9.44}$$

The displacement condition on $\Gamma_S$ is essential; the moment condition and both free-edge conditions are natural. Equations (9.43) and (9.44) are stated on smooth open edge segments. On polygonal boundaries, endpoint and corner contributions are understood through the global variational formulation rather than as additional pointwise edge equations.

Across an ideal panel seam, displacement and normal rotation would be continuous and the corresponding interface resultants would balance. In the present model, the two kinematic conditions are enforced weakly by Equation (9.21), while interface equilibrium follows from the variational statement rather than being imposed as an additional strong condition.

The initial data are

$$w(\cdot, 0; \boldsymbol{\mu}) = w_0, \qquad \dot{w}(\cdot, 0; \boldsymbol{\mu}) = v_0. \tag{9.45}$$

The reported snapshot campaigns use

$$w_0 = 0, \qquad v_0 = 0. \tag{9.46}$$

For the step forcing in Equation (9.27), these conditions do not imply zero initial acceleration; the latter is determined by the dynamic equilibrium at $\tau = 0$.

The Navier verification uses

$$\Gamma_S = \partial\Omega, \tag{9.47}$$

whereas the transient ROM campaigns use free exterior edges,

$$\Gamma_F = \partial\Omega. \tag{9.48}$$

## 9.7 Parameterized transient solution map

The discrete foundation-regime selection in Equation (9.18) is fixed when a campaign is defined; it is not an additional learned coordinate in the reported ROMs. Within a selected regime, collect the continuous physical and loading data in a parameter vector

$$\boldsymbol{\mu} = \left(D_x, D_y, D_{xy}, D_s, k_s, \varepsilon_f, \rho, t, A_q, \ldots\right) \in \mathcal{D} \subset \mathbb{R}^P. \tag{9.49}$$

Only a selected subset is varied in any one campaign. The interface penalties $\eta_w$ and $\eta_\psi$ are discretization data determined facetwise in Section 10.2; they are not sampled physical trajectory coordinates in the reported databases. A fixed value of $\boldsymbol{\mu}$ defines one initial-value problem over $[0, T]$. Time is the independent evolution variable and is not a material, geometric, or loading parameter. Nevertheless, the non-intrusive reduced models approximate the state at a specified parameter–time pair and therefore use the extended coordinate

$$\boldsymbol{\xi} = (\boldsymbol{\mu}, \tau) \in \mathcal{D} \times [0, T]. \tag{9.50}$$

For the active monolithic load-amplitude campaign,

$$\boldsymbol{\mu} = (A_q), \qquad \boldsymbol{\xi} = (A_q, \tau), \tag{9.51}$$

while the remaining physical quantities are fixed at their baseline values. The time horizon is

$$T = 0.20\,\mathrm{s}. \tag{9.52}$$

After spatial and temporal approximation, the high-fidelity displacement map is written as

$$\mathcal{S}_h : (\boldsymbol{\mu}, \tau) \longmapsto \widehat{w}_h(\boldsymbol{\mu}, \tau) \in X_h, \tag{9.53}$$

where $X_h$ is the common continuous comparison space used for snapshot storage, POD construction, and field-error evaluation. For a monolithic plate, $\widehat{w}_h$ is the finite element displacement itself. For a panelized configuration, it is the projected continuous representation of the panel-wise solution; that transfer is defined in the reduced-order section.

**Remark 9.7.1** (Scope of the transient model) For each prescribed effective-foundation regime, the transient model is linear and describes small transverse deflections with transverse translational inertia. The foundation regime is fixed throughout a trajectory, and the reduced-order models inherit the same kinematic, loading, and damping assumptions.

## 9.8 Summary and handoff to the full-order approximation

The equations above define a linear second-order evolution problem for each fixed physical parameter vector and prescribed support regime, within the scope stated in Section 9.1. In the later non-intrusive reduction, $(\boldsymbol{\mu}, \tau)$ is used as a query pair for the displacement field; $\tau$ remains the physical evolution variable, and the surrogate is not a causal reduced Newmark propagator.

Chapter 10 discretizes the fourth-order spatial operator with a panel-aligned symmetric $C^0$ interior-penalty method, constructs the active mass–stiffness system, advances it with the average-acceleration Newmark scheme, and verifies the implementation against complementary analytical and numerical references.

# Full-Order Approximation and Verification of the Transient Plate Model

# 10


10.1 Role and scope of the transient full-order model 111
10.2 Panel-aligned spatial discretization . . . . . 111
10.3 Semidiscrete dynamics and active degrees of freedom . . . . . 114
10.4 Time integration and snapshot storage . . . . . 115
10.5 Analytical and modal references . . . . . 116
10.6 Verification framework and results . . . . . 120
10.7 Summary and handoff to the reduced-order model 126


## 10.1 Role and scope of the transient full-order model

Chapter 9 defines a linear initial–boundary-value problem for each prescribed effective-foundation regime. This chapter constructs its high-fidelity finite element approximation. The fourth-order Kirchhoff–Love operator is treated by a panel-aligned symmetric $C^0$ interior-penalty formulation, while the dynamic setting additionally requires a consistent mass matrix, optional algebraic damping, admissible initial data, and a stable second-order time integrator.

The panel-partitioned realization introduces a specific algebraic issue. Each panel component is stored on the complete common mesh even though it carries physical mass only on its own subdomain. The assembled mixed vector therefore contains zero-mass component degrees of freedom. A single active-index restriction, determined from the mass diagonal, is applied consistently to the mass, stiffness, damping, load, initial data, generalized eigenproblem, direct and modal integration, state lifting, and snapshot generation. Retaining the inactive entries would produce a singular dynamic mass matrix and nonphysical algebraic states.

The verification studies supporting the transient snapshot databases are reported in Section 10.6 and summarized in Table 10.4. Their interpretation is restricted to the fixed-discretization scope stated in Section 10.6.

## 10.2 Panel-aligned spatial discretization

### Panel-wise finite element spaces and mixed representation

Let $\mathcal{T}_h$ be a conforming, shape-regular triangulation of $\Omega$ aligned with the physical panel partition, so that every interface $\Gamma_{ij}$ is a union of mesh facets. Its restriction to panel $\Omega_i$ is denoted by $\mathcal{T}_{h,i}$. For a polynomial degree $p \geq 2$, define the panel-wise continuous Lagrange space

$$W_{h,i} := \left\{ v_h \in C^0(\overline{\Omega_i}) : v_h|_K \in \mathbb{P}_p(K) \quad \forall K \in \mathcal{T}_{h,i} \right\}, \tag{10.1}$$

and the broken panel space

$$W_h := \bigoplus_{i=1}^{M} W_{h,i}. \tag{10.2}$$

The direct-sum notation expresses the physical approximation: fields belonging to adjacent panels retain independent traces before the interface terms are applied. The discrete admissible space is

$$V_h := \{ v_h \in W_h : v_h = 0 \text{ on } \Gamma_S \}. \tag{10.3}$$

Thus, the displacement condition on a simply supported edge is imposed strongly, whereas the corresponding moment condition and the free-edge

conditions remain natural. If $\Gamma_S = \varnothing$, no essential exterior condition is imposed.

In the implementation, the panel fields are stored as components of a mixed finite element function defined on the common mesh. The $i$-th component is physical only on $\Omega_i$, and the panel-wise field is interpreted as

$$w_h|_{\Omega_i} = w_{h,i}|_{\Omega_i}. \tag{10.4}$$

This global mixed realization is algebraically convenient, but it introduces inactive component degrees of freedom outside their physical panels. Their consistent treatment in the dynamic system is addressed in Section 10.3.

## Symmetric $C^0$ interior-penalty form

Let $\mathscr{E}_{h,i}^{\mathrm{int}}$ denote the set of mesh facets whose two adjacent elements both belong to $\Omega_i$. For $F = \partial K^+ \cap \partial K^- \in \mathscr{E}_{h,i}^{\mathrm{int}}$, let $\mathbf{n}^\pm$ be the outward unit normals of $K^\pm$, and let

$$h_F := \frac{1}{2}\,(h_{K^+} + h_{K^-})$$

be the arithmetic mean of the adjacent cell diameters. The orientation-invariant normal-slope jump and the arithmetic average are

$$[\![\partial_n v_h]\!]_F := \nabla v_h^+ \cdot \mathbf{n}^+ + \nabla v_h^- \cdot \mathbf{n}^-, \qquad \{\!\{\zeta\}\!\}_F := \frac{1}{2}(\zeta^+ + \zeta^-). \tag{10.5}$$

These traces are associated with numerical element facets and are distinct from the physical panel-interface jumps $[\![\cdot]\!]_{ij}$ introduced in Section 9.2.

For a scalar field $v_h$, introduce the curvature-conjugate moment tensor

$$\mathbf{B}(v_h) := \begin{bmatrix} D_x v_{h,xx} + D_{xy} v_{h,yy} & 2D_s v_{h,xy} \\ 2D_s v_{h,xy} & D_{xy} v_{h,xx} + D_y v_{h,yy} \end{bmatrix}, \tag{10.6}$$

and define its one-sided normal–normal traces by

$$M_{nn}^\pm(v_h) := (\mathbf{n}^\pm)^{\mathsf{T}} \mathbf{B}(v_h^\pm)\mathbf{n}^\pm. \tag{10.7}$$

Because the normal enters quadratically, $M_{nn}^\pm$ is unaffected by reversal of the facet orientation. With

$$D_{\mathrm{eff}} := \sqrt{D_x D_y}, \qquad \tau_F := (p+2)^2 \frac{D_{\mathrm{eff}}}{h_F}, \tag{10.8}$$

the symmetric $C^0$ interior-penalty bending form on $\Omega_i$ is

$$
\begin{aligned}
a^{b}_{h,i}(w_h, v_h) := \sum_{K\in\mathcal{T}_{h,i}} \int_K \Big[ & D_x w_{h,xx} v_{h,xx} + D_y w_{h,yy} v_{h,yy} \\
& + D_{xy}\left(w_{h,xx} v_{h,yy} + w_{h,yy} v_{h,xx}\right) \\
& + 4 D_s w_{h,xy} v_{h,xy} \Big] \, \mathrm{d}\Omega \\
- \sum_{F\in\mathcal{E}^{\mathrm{int}}_{h,i}} \int_F & \{\!\{ M_{nn}(w_h) \}\!\}_F \, [\![ \partial_n v_h ]\!]_F \, \mathrm{d}\Gamma \\
- \sum_{F\in\mathcal{E}^{\mathrm{int}}_{h,i}} \int_F & \{\!\{ M_{nn}(v_h) \}\!\}_F \, [\![ \partial_n w_h ]\!]_F \, \mathrm{d}\Gamma \\
+ \sum_{F\in\mathcal{E}^{\mathrm{int}}_{h,i}} \int_F & \tau_F \, [\![ \partial_n w_h ]\!]_F \, [\![ \partial_n v_h ]\!]_F \, \mathrm{d}\Gamma .
\end{aligned}
\tag{10.9}
$$

The volume term represents orthotropic bending, the next two terms provide consistency and symmetry, and the last term weakly controls the discontinuous normal slope. This is the only continuity missing from the panel-wise $C^0$ space on ordinary element facets [25–27, 38].

For a mixed multi-panel realization, the $i$-th component is globally defined algebraically even though it is physical only on $\Omega_i$. The implemented facet contribution is therefore restricted to facets whose two adjacent cells belong to the same physical panel. This restriction is selected through piecewise-constant panel indicator functions. It does not rely on the arbitrary plus/minus orientation assigned by the finite element library and prevents a physical seam from being treated as an ordinary inter-element facet.

## Physical seam penalties, stiffness, and mass

The physical panel-interface contribution is assembled separately from Equation (10.9). On each interface facet $F \subset \Gamma_{ij}$, the facetwise weights used in the reported computations are

$$
\eta_{w,F} := p \, \frac{D_{\mathrm{eff}}}{h_F^3}, \qquad \eta_{\psi,F} := \frac{12^{p+1}}{p} \, \frac{D_{\mathrm{eff}}}{h_F}. \tag{10.10}
$$

The dimensions are $\eta_{w,F} \sim \mathrm{N\,m^{-2}}$ and $\eta_{\psi,F} \sim \mathrm{N}$, so both interface contributions have the dimensions of energy after integration over a line facet. These weights define the interface scaling used in the reported computations; they are not asserted to be universal or parameter-optimal penalty choices.

The corresponding discrete seam form is

$$
\begin{aligned}
c_{\Gamma,h}(w_h, v_h) := \sum_{1\le i<j\le M} \sum_{F\subset\Gamma_{ij}} \int_F \Big[ & \eta_{w,F} \, [\![ w_h ]\!]_{ij} \, [\![ v_h ]\!]_{ij} \\
& + \eta_{\psi,F} \, [\![ \partial_n w_h ]\!]_{ij} \, [\![ \partial_n v_h ]\!]_{ij} \Big] \, \mathrm{d}\Gamma .
\end{aligned}
\tag{10.11}
$$

The physical traces in Equation (10.11) are selected by panel indicators, so the trace associated with $\Omega_i$ is identified by its physical support rather than by a plus/minus facet label.

The complete discrete stiffness form is

$$k_{h,\boldsymbol{\mu}}(w_h, v_h) := \sum_{i=1}^{M} a^b_{h,i}(w_h, v_h) + \sum_{i=1}^{M} \int_{\Omega_i} k_s^{\mathrm{eff}} w_{h,i} v_{h,i} \, \mathrm{d}\Omega + c_{\Gamma,h}(w_h, v_h), \tag{10.12}$$

and the consistent transverse mass form is

$$m_h(w_h, v_h) := \sum_{i=1}^{M} \int_{\Omega_i} \rho t \, w_{h,i} v_{h,i} \, \mathrm{d}\Omega. \tag{10.13}$$

The retained operator is the consistent Galerkin mass rather than a diagonal mass-lumped approximation [97, 98]. The construction separates four conceptually different ingredients: $C^0$-IPG terms act only on numerical facets inside a panel; seam penalties act only on physical interfaces; boundary conditions act on $\partial\Omega$; and the foundation contributes a distributed zero-order stiffness.

## 10.3 Semidiscrete dynamics and active degrees of freedom

Let $\{\varphi_j\}_{j=1}^{N_{\mathrm{mix}}}$ be a basis of the assembled mixed space and write

$$w_h(\tau; \boldsymbol{\mu}) = \sum_{j=1}^{N_{\mathrm{mix}}} u_j(\tau; \boldsymbol{\mu}) \varphi_j. \tag{10.14}$$

Before essential-boundary elimination, the assembled matrices satisfy

$$[\widetilde{\mathbf{M}}]_{ij} = m_h(\varphi_j, \varphi_i), \qquad [\widetilde{\mathbf{K}}]_{ij} = k_{h,\boldsymbol{\mu}}(\varphi_j, \varphi_i). \tag{10.15}$$

Let $\mathcal{J}_{\mathrm{free}}$ denote the indices remaining after the strong boundary conditions have been removed. In a monolithic realization, all remaining degrees of freedom possess physical mass. In the global mixed realization of a panel partition, however, a component basis function may lie outside that component's physical panel. Its mass contribution then vanishes, and retaining it would render the dynamic mass matrix singular without adding a physical state.

The implementation identifies the dynamically active indices from the mass diagonal. With

$$\tau_M := 10^{-12} \max_{k \in \mathcal{J}_{\mathrm{free}}} \left| \widetilde{M}_{kk} \right|, \tag{10.16}$$

the active set is

$$\mathcal{J}_{\mathrm{dyn}} := \left\{ j \in \mathcal{J}_{\mathrm{free}} : \left| \widetilde{M}_{jj} \right| > \tau_M \right\}. \tag{10.17}$$

Exactly the same index set is used for the active mass, stiffness, and damping matrices; the generalized eigenproblem; direct and modal Newmark integration; restriction of the load and initial-condition vectors; active-state lifting; reconstruction of the physical panel-wise and global output fields; and snapshot storage. This filtering is an algebraic consequence of the global mixed representation and does not modify the physical panel model.

On the active space, the semidiscrete equations are

$$\mathbf{M}(\boldsymbol{\mu})\ddot{\mathbf{u}}(\tau;\boldsymbol{\mu}) + \mathbf{C}(\boldsymbol{\mu})\dot{\mathbf{u}}(\tau;\boldsymbol{\mu}) + \mathbf{K}(\boldsymbol{\mu})\mathbf{u}(\tau;\boldsymbol{\mu}) = \mathbf{f}(\tau;\boldsymbol{\mu}), \tag{10.18}$$

where $\mathbf{M}$ and $\mathbf{K}$ are the active restrictions of $\widetilde{\mathbf{M}}$ and $\widetilde{\mathbf{K}}$. Optional damping is introduced algebraically through

$$\mathbf{C}(\boldsymbol{\mu}) = \alpha_M \mathbf{M}(\boldsymbol{\mu}) + \alpha_K \mathbf{K}(\boldsymbol{\mu}), \qquad [\alpha_M] = \mathrm{s}^{-1}, \quad [\alpha_K] = \mathrm{s}. \tag{10.19}$$

The active snapshot campaigns use the zero values reported in Equation (9.37); the harmonic verification below instead introduces a modal damping ratio and does not use Equation (10.19). The load vector is assembled from the spatial load and multiplied by the prescribed temporal factor. The vectors $\mathbf{u}_0$ and $\mathbf{v}_0$ are the active-degree-of-freedom representations of the initial displacement and velocity.

## 10.4 Time integration and snapshot storage

Let $\tau_n = n\Delta\tau$, $n = 0, \ldots, N_{\mathrm{step}}$, with a uniform step $\Delta\tau = T/N_{\mathrm{step}}$. Denote the discrete displacement, velocity, and acceleration by $\mathbf{u}_n$, $\mathbf{v}_n$, and $\mathbf{a}_n$. Newmark's relations are

$$\mathbf{u}_{n+1} = \mathbf{u}_n + \Delta\tau\mathbf{v}_n + \Delta\tau^2\left[\left(\frac{1}{2} - \beta\right)\mathbf{a}_n + \beta\mathbf{a}_{n+1}\right], \tag{10.20}$$

$$\mathbf{v}_{n+1} = \mathbf{v}_n + \Delta\tau\left[(1-\gamma)\mathbf{a}_n + \gamma\mathbf{a}_{n+1}\right]. \tag{10.21}$$

The solver uses the average-acceleration parameters

$$\beta = \frac{1}{4}, \qquad \gamma = \frac{1}{2}. \tag{10.22}$$

For a linear semidiscrete system with symmetric positive-definite mass and symmetric positive-semidefinite stiffness and damping, this member of the Newmark family is second-order accurate and unconditionally stable in the standard linear sense [13, 34, 95, 99]. This statement concerns the fixed-regime linear system in Equation (10.18); it does not extend to an unimplemented state-switching contact problem.

Introduce

$$\begin{aligned} a_0 &= \frac{1}{\beta\Delta\tau^2}, \quad a_1 = \frac{\gamma}{\beta\Delta\tau}, \quad a_2 = \frac{1}{\beta\Delta\tau}, \\ a_3 &= \frac{1}{2\beta} - 1, \quad a_4 = \frac{\gamma}{\beta} - 1, \quad a_5 = \Delta\tau\left(\frac{\gamma}{2\beta} - 1\right). \end{aligned} \tag{10.23}$$

Since the effective foundation regime and the Rayleigh coefficients are fixed within a trajectory, the effective matrix

$$\mathbf{K}_{\mathrm{eff}} := \mathbf{K} + a_0\mathbf{M} + a_1\mathbf{C} \tag{10.24}$$

is constant over the time interval. At each step, the displacement is obtained from

$$\mathbf{K}_{\mathrm{eff}}\mathbf{u}_{n+1} = \mathbf{f}_{n+1} + \mathbf{M}\left(a_0\mathbf{u}_n + a_2\mathbf{v}_n + a_3\mathbf{a}_n\right) + \mathbf{C}\left(a_1\mathbf{u}_n + a_4\mathbf{v}_n + a_5\mathbf{a}_n\right), \tag{10.25}$$

followed by

$$\mathbf{a}_{n+1} = a_0(\mathbf{u}_{n+1} - \mathbf{u}_n) - a_2\mathbf{v}_n - a_3\mathbf{a}_n, \tag{10.26}$$

$$\mathbf{v}_{n+1} = \mathbf{v}_n + \Delta\tau\,[(1-\gamma)\mathbf{a}_n + \gamma\mathbf{a}_{n+1}]\,. \tag{10.27}$$

Here $\mathbf{f}_{n+1} := \mathbf{f}(\tau_{n+1};\boldsymbol{\mu})$; thus, the load is evaluated at the end of the step. The initial acceleration is determined by

$$\mathbf{M}\mathbf{a}_0 = \mathbf{f}(\tau_0) - \mathbf{C}\mathbf{v}_0 - \mathbf{K}\mathbf{u}_0. \tag{10.28}$$

For the implemented step profile, $\mathbf{f}(\tau_0) \neq \mathbf{0}$ because the load factor is set to one at $\tau = 0$.

The high-resolution snapshot campaigns use

$$T = 0.20\,\mathrm{s}, \qquad \Delta\tau = 10^{-3}\,\mathrm{s}, \qquad N_{\mathrm{step}} = 200, \tag{10.29}$$

with storage every five integration steps. Hence,

$$\left\{\tau_r^{\mathrm{store}}\right\}_{r=0}^{40} = \{0, 0.005, 0.010, \ldots, 0.200\}\ \mathrm{s}, \qquad N_\tau = 41. \tag{10.30}$$

Those campaigns use the integer mesh-resolution selector $N_\Omega = 64$ and quadratic Lagrange elements. In the implementation, $N_\Omega$ is the resolution argument supplied to the unstructured mesh generator; it is not the number of cells in either coordinate direction. The verification calculations below use the same time step but $N_\Omega = 32$. The reported finite element dimensions, rather than $N_\Omega$ alone, provide the unambiguous algebraic sizes of the reported systems.

## 10.5 Analytical and modal references

### Finite element eigenproblem and Navier frequencies

For free vibration, let $\mathbf{u}(\tau) = \boldsymbol{\phi} e^{\mathrm{i}\omega\tau}$. The active finite element matrices satisfy the generalized eigenproblem

$$\mathbf{K}\boldsymbol{\phi}_j = \omega_j^2 \mathbf{M}\boldsymbol{\phi}_j, \qquad \boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{M}\boldsymbol{\phi}_j = \delta_{ij}. \tag{10.31}$$

The generalized problem and its mass normalization follow standard structural modal-analysis and eigen-solution practice [51, 95, 100]. For a monolithic rectangle with constant coefficients and all four edges simply supported, the Navier modes are

$$W_{mn}(x,y) = \sin(\alpha_m x)\sin(\beta_n y), \qquad \alpha_m = \frac{m\pi}{L}, \qquad \beta_n = \frac{n\pi}{B}. \tag{10.32}$$

With the transverse-only inertia retained in the present model, substitution into the strong equation gives

$$\omega_{mn}^2 = \frac{D_x\alpha_m^4 + 2(D_{xy} + 2D_s)\alpha_m^2\beta_n^2 + D_y\beta_n^4 + k_s^{\mathrm{eff}}}{\rho t}. \tag{10.33}$$

Reddy's general classical-plate dynamic equation contains the translational inertia $I_0 = \rho t$ and may also retain the rotary-inertia coefficient $I_2 = \rho t^3/12$. The model used here sets $I_2 = 0$; accordingly, Equation (10.33) contains no rotary-inertia correction [13].

The verification study uses two references whose spatial structure can be worked out analytically. The first is the Navier frequency formula for a simply supported aligned-orthotropic rectangle with a constant prescribed foundation coefficient. The second is a single Navier mode forced in time and advanced with the same average-acceleration Newmark scheme as the finite element system. Including both derivations here makes the normalization, inertia, load projection, and temporal comparison reproducible without relying on an external appendix.

## Navier frequency derivation for an aligned-orthotropic plate

The derivation applies to a monolithic, aligned-orthotropic rectangular plate with constant coefficients, a prescribed constant effective Winkler coefficient, and simply supported conditions on all four edges. It therefore provides a reference for the selected linear foundation regime rather than for a state-switching unilateral-contact problem.

Let

$$\Omega = (0, L) \times (0, B), \tag{10.34}$$

with constant rigidities, mass density, thickness, and effective foundation coefficient $k_s^{\text{eff}}$. In the absence of damping, the strong transverse-equilibrium equation is

$$\begin{aligned} \rho t\, \partial_{\tau\tau} w + D_x \partial_{xxxx} w + 2(D_{xy} + 2D_s)\partial_{xxyy} w \\ + D_y \partial_{yyyy} w + k_s^{\text{eff}} w = q \qquad \text{in } \Omega \times (0, T). \end{aligned} \tag{10.35}$$

For free vibration, set $q = 0$ and seek a separated solution

$$w_{mn}(x, y, \tau) = \Phi_{mn}(x, y) Q_{mn}(\tau), \qquad \Phi_{mn}(x, y) = \sin(\alpha_m x) \sin(\beta_n y), \tag{10.36}$$

where

$$\alpha_m = \frac{m\pi}{L}, \qquad \beta_n = \frac{n\pi}{B}, \qquad m, n = 1, 2, \ldots. \tag{10.37}$$

The sine field vanishes on every exterior edge. On $x = 0, L$,

$$\Phi_{mn,xx} = -\alpha_m^2 \Phi_{mn} = 0, \qquad \Phi_{mn,yy} = -\beta_n^2 \Phi_{mn} = 0, \tag{10.38}$$

and hence

$$M_{nn}(\Phi_{mn}) = M_x(\Phi_{mn}) = D_x \Phi_{mn,xx} + D_{xy} \Phi_{mn,yy} = 0. \tag{10.39}$$

The analogous argument on $y = 0, B$ gives

$$M_{nn}(\Phi_{mn}) = M_y(\Phi_{mn}) = D_{xy} \Phi_{mn,xx} + D_y \Phi_{mn,yy} = 0. \tag{10.40}$$

Thus, $\Phi_{mn}$ satisfies both the essential displacement condition and the natural normal-moment condition on all four simply supported edges.

The required spatial derivatives are

$$\begin{aligned} &\Phi_{mn,xxxx} = \alpha_m^4 \Phi_{mn}, \qquad \Phi_{mn,xxyy} = \alpha_m^2 \beta_n^2 \Phi_{mn}, \\ &\Phi_{mn,yyyy} = \beta_n^4 \Phi_{mn}. \end{aligned} \tag{10.41}$$

Substitution of Equations (10.36) and (10.41) into Equation (10.35) yields

$$\rho t\, \ddot{Q}_{mn} + \left[ D_x \alpha_m^4 + 2(D_{xy} + 2D_s)\alpha_m^2 \beta_n^2 + D_y \beta_n^4 + k_s^{\text{eff}} \right] Q_{mn} = 0. \tag{10.42}$$

For a harmonic modal coordinate

$$Q_{mn}(\tau) = \widehat{Q}_{mn} \exp(\mathrm{i}\omega_{mn}\tau), \qquad \widehat{Q}_{mn} \neq 0,$$

the natural angular frequency satisfies

$$\omega_{mn}^2 = \frac{D_x\alpha_m^4 + 2(D_{xy} + 2D_s)\alpha_m^2\beta_n^2 + D_y\beta_n^4 + k_s^{\mathrm{eff}}}{\rho t}. \tag{10.43}$$

The denominator is the retained transverse mass per unit area. Consistently with the implemented Kirchhoff–Love model, Equation (10.43) contains neither a rotary-inertia correction nor a transverse-shear contribution. The formula also assumes that $k_s^{\mathrm{eff}}$ is prescribed and spatially constant during the modal calculation.

## Spatially analytical single-mode transient reference

The spatial reduction is analytical, whereas the comparison trajectory is advanced with the same Newmark scheme and time grid as the finite element system. The construction therefore isolates the finite element spatial operator, modal normalization, load assembly, and direct time-integration implementation without introducing a mismatch in the temporal discretization.

Let the simply supported plate be forced in one Navier mode,

$$q(x, y, \tau) = q_0\Phi_{mn}(x, y)s(\tau), \tag{10.44}$$

where $q_0 \in \mathbb{R}$ follows the upward-positive load convention and $\Phi_{mn}$ is defined in Equation (10.36). Assume a response of the form

$$w(x, y, \tau) = \Phi_{mn}(x, y)Q_{mn}(\tau). \tag{10.45}$$

Multiplication of Equation (10.35) by $\Phi_{mn}$ and integration over Equation (10.34) reduce the plate equation to a scalar modal problem. Since

$$\begin{aligned}\int_\Omega \Phi_{mn}^2 \,\mathrm{d}\Omega &= \left(\int_0^L \sin^2(\alpha_m x)\,\mathrm{d}x\right)\left(\int_0^B \sin^2(\beta_n y)\,\mathrm{d}y\right) \\ &= \frac{LB}{4},\end{aligned} \tag{10.46}$$

the unnormalized modal mass, stiffness, and load are

$$m_{mn} = \rho t\frac{LB}{4}, \tag{10.47}$$

$$k_{mn} = m_{mn}\omega_{mn}^2, \tag{10.48}$$

$$p_{mn}(\tau) = q_0\frac{LB}{4}s(\tau). \tag{10.49}$$

The generalized coordinate satisfies

$$m_{mn}\ddot{Q}_{mn} + c_{mn}\dot{Q}_{mn} + k_{mn}Q_{mn} = p_{mn}(\tau). \tag{10.50}$$

If a modal damping ratio $\zeta_{mn}$ is prescribed, the corresponding scalar damping coefficient is

$$c_{mn} = 2\zeta_{mn}\omega_{mn}m_{mn}. \tag{10.51}$$

The direct single-mode transient verification is undamped, so that $c_{mn} = 0$ in the reported comparison.

For the implemented step factor $s(\tau) = H(\tau)$, with $H(0) = 1$, define

$$p_{mn}^{0} = q_0 \frac{LB}{4}. \tag{10.52}$$

With

$$Q_{mn}(0) = 0, \qquad \dot{Q}_{mn}(0) = 0, \tag{10.53}$$

the continuous-time undamped response is

$$Q_{mn}(\tau) = \frac{p_{mn}^{0}}{k_{mn}} \left[1 - \cos(\omega_{mn}\tau)\right], \qquad \tau \geq 0. \tag{10.54}$$

Because the load is active at $\tau = 0$, the initial modal acceleration is nonzero:

$$\ddot{Q}_{mn}(0) = \frac{p_{mn}^{0}}{m_{mn}}. \tag{10.55}$$

Equation (10.54) is retained as a continuous-time consistency relation, but it is not the reference history used in the reported finite element comparison.

To define the time-discrete reference, let

$$\tau_r = r\Delta\tau, \qquad r = 0, \ldots, N_{\text{step}}, \tag{10.56}$$

and denote the modal displacement, velocity, and acceleration by $Q_{mn}^{r}$, $\dot{Q}_{mn}^{r}$, and $\ddot{Q}_{mn}^{r}$, respectively. Using the Newmark coefficients $a_0, \ldots, a_5$ introduced in Equation (10.23), the scalar effective stiffness is

$$\begin{aligned} k_{mn}^{\text{eff}} &= k_{mn} + a_0 m_{mn} + a_1 c_{mn} \\ &= k_{mn} + \frac{1}{\beta\Delta\tau^2} m_{mn} + \frac{\gamma}{\beta\Delta\tau} c_{mn}, \qquad \beta = \frac{1}{4}, \qquad \gamma = \frac{1}{2}. \end{aligned} \tag{10.57}$$

At time $\tau_{r+1}$, the modal displacement is obtained from

$$\begin{aligned} k_{mn}^{\text{eff}} Q_{mn}^{r+1} = {} & p_{mn}(\tau_{r+1}) + m_{mn}\left(a_0 Q_{mn}^{r} + a_2 \dot{Q}_{mn}^{r} + a_3 \ddot{Q}_{mn}^{r}\right) \\ & + c_{mn}\left(a_1 Q_{mn}^{r} + a_4 \dot{Q}_{mn}^{r} + a_5 \ddot{Q}_{mn}^{r}\right). \end{aligned} \tag{10.58}$$

The acceleration and velocity are then updated through

$$\ddot{Q}_{mn}^{r+1} = a_0\left(Q_{mn}^{r+1} - Q_{mn}^{r}\right) - a_2 \dot{Q}_{mn}^{r} - a_3 \ddot{Q}_{mn}^{r}, \tag{10.59}$$

$$\dot{Q}_{mn}^{r+1} = \dot{Q}_{mn}^{r} + \Delta\tau\left[(1-\gamma)\ddot{Q}_{mn}^{r} + \gamma \ddot{Q}_{mn}^{r+1}\right]. \tag{10.60}$$

The initial acceleration is determined from

$$m_{mn}\ddot{Q}_{mn}^{0} = p_{mn}(\tau_0) - c_{mn}\dot{Q}_{mn}^{0} - k_{mn}Q_{mn}^{0}. \tag{10.61}$$

Equations (10.58)–(10.60) are the scalar counterparts of Equations (10.25), (10.26), and (10.27). The modal load is evaluated at $\tau_{r+1}$, exactly as in the direct finite element trajectory. The resulting reference is therefore analytical in space and discrete in time, with the same average-acceleration Newmark parameters and time grid as the full-order comparison.

## 10.6 Verification framework and results

The verification program combines seven complementary component-level checks. Every test uses a monolithic or panelized $2\,\mathrm{m} \times 1\,\mathrm{m}$ plate, quadratic elements, the integer mesh-resolution selector $N_\Omega = 32$, and simply supported exterior edges. The time-domain tests use $T = 0.20\,\mathrm{s}$ and $\Delta\tau = 10^{-3}\,\mathrm{s}$; the generalized eigenvalue and harmonic-response checks use the same spatial discretization but do not advance a trajectory on that time grid.

The monolithic single-mode, modal, harmonic, and patch-load histories are observed at the center point

$$P_c = (1.0, 0.5)\ \mathrm{m}, \tag{10.62}$$

whereas the panel-to-monolithic comparison uses the off-center point

$$P_p = (0.73, 0.37)\ \mathrm{m}. \tag{10.63}$$

The latter avoids placing the comparison probe on the vertical panel interface. The numerical material preset is an isotropic specialization of the aligned-orthotropic formulation:

$$\begin{aligned} E_x &= E_y = E, & G_{xy} &= \frac{E}{2(1+\nu_{xy})}, \\ D_x &= D_y = \frac{Et^3}{12(1-\nu_{xy}^2)}, & & \\ D_{xy} &= \nu_{xy} D_x, & D_s &= \frac{1-\nu_{xy}}{2} D_x. \end{aligned} \tag{10.64}$$

The resulting values are reported in Table 10.1. This preset verifies the implemented plate operator in an isotropic limiting case; it does not constitute a separate anisotropic-parameter verification.

**Table 10.1:** Isotropic material specialization used in the full-order verification program.

| Quantity | Value | Unit |
|---|---|---|
| $E_x = E_y$ | $9.0 \times 10^8$ | Pa |
| $\nu_{xy}$ | 0.30 | – |
| $G_{xy}$ | $3.461538 \times 10^8$ | Pa |
| $\rho$ | $1.2 \times 10^3$ | $\mathrm{kg\,m^{-3}}$ |
| $t$ | $5.0 \times 10^{-2}$ | m |
| $k_s$ | $1.0 \times 10^5$ | $\mathrm{N\,m^{-3}}$ |
| $D_x = D_y$ | $1.030220 \times 10^4$ | N m |
| $D_{xy}$ | $3.090659 \times 10^3$ | N m |
| $D_s$ | $3.605769 \times 10^3$ | N m |

For two histories sampled on the same time grid, let $y_n^{\mathrm{ref}}$ denote the reference history and $y_n^{\mathrm{cmp}}$ the comparison history. The reported relative discrepancies are

$$e_{\ell^2} := \frac{\left(\sum_n \left|y_n^{\mathrm{cmp}} - y_n^{\mathrm{ref}}\right|^2\right)^{1/2}}{\left(\sum_n \left|y_n^{\mathrm{ref}}\right|^2\right)^{1/2}}, \qquad e_{\ell^\infty} := \frac{\max_n \left|y_n^{\mathrm{cmp}} - y_n^{\mathrm{ref}}\right|}{\max_n \left|y_n^{\mathrm{ref}}\right|}. \tag{10.65}$$

For natural frequencies,

$$e_{\omega,j} := \frac{\left|\omega_{h,j} - \omega_j\right|}{\omega_j}. \tag{10.66}$$

## Verification results

**Natural frequencies against the Navier formula.** The first eight finite element frequencies are compared with the lowest eight ordered Navier frequencies for three prescribed effective-foundation regimes: no foundation, compression, and regularized uplift with $\varepsilon_f = 10^{-3}$. The active simply supported finite element system contains 4,908 dynamic degrees of freedom. The regime-specific results are reported in Table 10.2. Across all tested modes and branches, the largest relative frequency discrepancy is

$$\max e_\omega = 7.346886 \times 10^{-3}. \tag{10.67}$$

Within the isotropic specialization in Equation (10.64), this comparison simultaneously checks the assembled fourth-order bending operator, the effective foundation contribution, the consistent transverse mass matrix, the generalized eigenvalue solve, and the simply supported boundary treatment.

**Table 10.2:** Finite element natural-frequency verification against the Navier formula.

| Regime | $k_s^{\mathrm{eff}}$ [N m$^{-3}$] | First $f_h$ [Hz] | Max. $e_\omega$ | Mean $e_\omega$ |
|---|---|---|---|---|
| No foundation | 0 | 25.75975 | $7.346886 \times 10^{-3}$ | $4.263492 \times 10^{-3}$ |
| Compression | $1.0 \times 10^5$ | 26.56656 | $7.333034 \times 10^{-3}$ | $4.233219 \times 10^{-3}$ |
| Regularized uplift | $1.0 \times 10^2$ | 25.76057 | $7.346872 \times 10^{-3}$ | $4.263461 \times 10^{-3}$ |

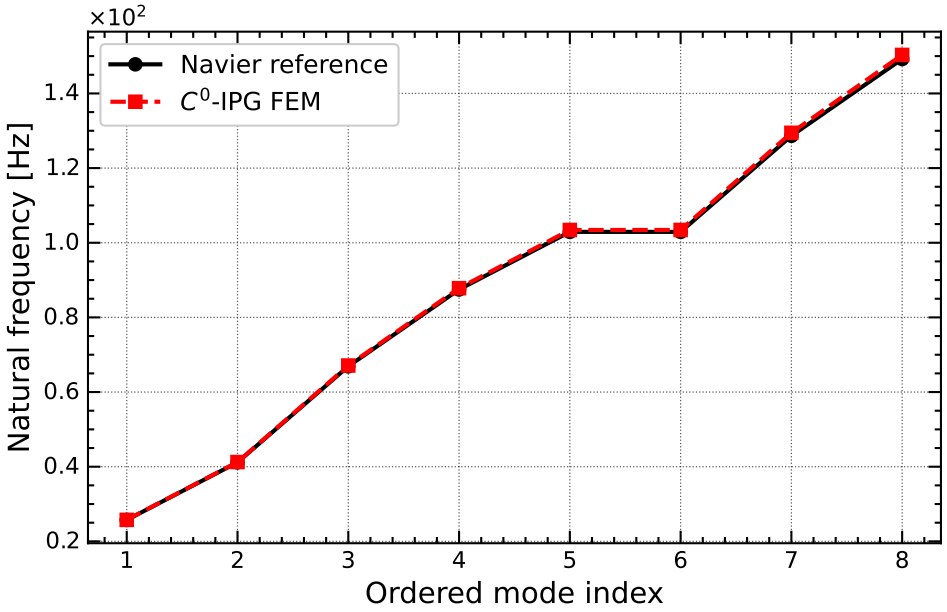



**Figure 10.1:** First eight natural frequencies for the foundation-free, simply supported benchmark at $N_\Omega = 32$ and $p = 2$, compared with the ordered Navier frequencies.

For completeness, the no-foundation values are listed in Table 10.3. The discrepancy increases gradually with mode rank at the fixed mesh, as expected for a spatially discretized eigenvalue problem.

**Table 10.3:** First eight foundation-free frequencies in the simply supported benchmark. Displayed frequencies are rounded; the relative errors are computed from the unrounded angular frequencies.

| Ordered mode | Navier [Hz] | FEM [Hz] | Relative error |
|---|---|---|---|
| 1 | 25.72878 | 25.7598 | $1.203737 \times 10^{-3}$ |
| 2 | 41.16605 | 41.2476 | $1.981510 \times 10^{-3}$ |
| 3 | 66.89483 | 67.1140 | $3.276628 \times 10^{-3}$ |
| 4 | 87.47785 | 87.8314 | $4.041058 \times 10^{-3}$ |
| 5 | 102.9151 | 103.4225 | $4.930243 \times 10^{-3}$ |
| 6 | 102.9151 | 103.4306 | $5.008939 \times 10^{-3}$ |
| 7 | 128.6439 | 129.4568 | $6.318939 \times 10^{-3}$ |
| 8 | 149.2269 | 150.3233 | $7.346886 \times 10^{-3}$ |

**Direct Newmark response against a Navier single-mode reference.** A spatial load proportional to the first Navier mode,

$$q(x, y, \tau) = q_0 \sin\left(\frac{\pi x}{L}\right) \sin\left(\frac{\pi y}{B}\right) H(\tau), \qquad q_0 = 1000\,\mathrm{N\,m^{-2}}, \quad (10.68)$$

is used to isolate the first modal response. The reference is analytical in space: orthogonality reduces the continuous plate problem to one scalar modal equation. That scalar equation is then advanced with the same Newmark scheme and time grid as the finite element system. The comparison is therefore spatially analytical and temporally discrete rather than an exact continuous-time benchmark. It checks the load normalization, transverse mass scaling, simply supported boundary treatment, and direct Newmark implementation.

The relative $\ell^2$ errors are $1.312629 \times 10^{-2}$, $1.314488 \times 10^{-2}$, and $1.312619 \times 10^{-2}$ for the no-foundation, compression, and regularized uplift regimes, respectively. The largest relative $\ell^\infty$ error over the three cases is $1.906461 \times 10^{-2}$.

**Direct and finite element modal Newmark integration.** The same semidiscrete systems are projected onto their first eight mass-normalized finite element eigenmodes and integrated with Newmark's method. Because the direct and modal calculations use the same assembled matrices, first-mode forcing, and time grid, this comparison isolates modal projection, load reduction, and field reconstruction. The very small discrepancies are expected for this deliberately single-mode-dominated test and must not be interpreted as an eight-mode accuracy estimate for arbitrary loading. The largest direct-to-modal relative discrepancies over the three effective foundation regimes are

$$\max e_{\ell^2} = 8.470401 \times 10^{-7}, \qquad \max e_{\ell^\infty} = 9.458344 \times 10^{-7}. \quad (10.69)$$

**Undamped energy conservation.** The unforced, foundation-free, undamped, simply supported plate is initialized with

$$w_h(x, y, 0) = 10^{-4} \sin\left(\frac{\pi x}{L}\right) \sin\left(\frac{\pi y}{B}\right), \qquad \dot{w}_h(x, y, 0) = 0. \quad (10.70)$$

For the active finite element system, the discrete mechanical energy is

$$E_n := \frac{1}{2} \mathbf{v}_n^\mathsf{T} \mathbf{M} \mathbf{v}_n + \frac{1}{2} \mathbf{u}_n^\mathsf{T} \mathbf{K} \mathbf{u}_n. \quad (10.71)$$

**(a) Direct versus Navier single-mode reference.** **(b) Direct versus eight-mode FE response.**

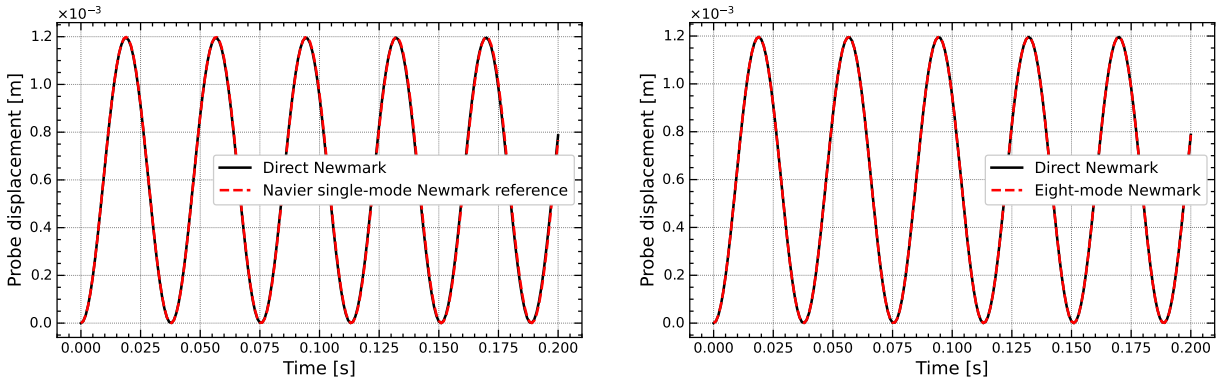



**Figure 10.2:** Time-domain verification at the center probe for the prescribed compression coefficient, $N_\Omega = 32$, $p = 2$, and $\Delta\tau = 10^{-3}$ s. Panel (a) compares direct Newmark integration with the spatially analytical, time-discrete single-mode reference. Panel (b) compares direct and eight-mode finite element Newmark responses under the same first-mode forcing. The discrepancies for all three effective-foundation regimes are reported in the text.

The computed history gives

$$\frac{\max_n E_n - \min_n E_n}{\operatorname{mean}_n |E_n|} = 1.039256 \times 10^{-10}. \tag{10.72}$$

This result verifies the algebraic consistency of the undamped average-acceleration update for the tested linear system [101].

**Harmonic frequency-response check.** Under the prescribed compression regime, eight mass-normalized finite element modes are used in a harmonic-response calculation. A uniform modal damping ratio $\zeta = 0.02$ is introduced for this frequency-domain diagnostic only; it is separate from the zero-Rayleigh-damping setting of the sampled transient campaigns. The excitation angular frequency is sampled at 180 points over

$$\frac{\omega_{\mathrm{exc}}}{\omega_1} \in [0.20, 2.50].$$

The largest sampled probe amplitude occurs at

$$f_{\mathrm{peak}} = 26.477507\,\mathrm{Hz}, \qquad f_1 = 26.566557\,\mathrm{Hz}, \tag{10.73}$$

so that

$$\frac{f_{\mathrm{peak}}}{f_1} = 0.996648, \qquad \frac{|f_{\mathrm{peak}} - f_1|}{f_1} = 3.351955 \times 10^{-3}. \tag{10.74}$$

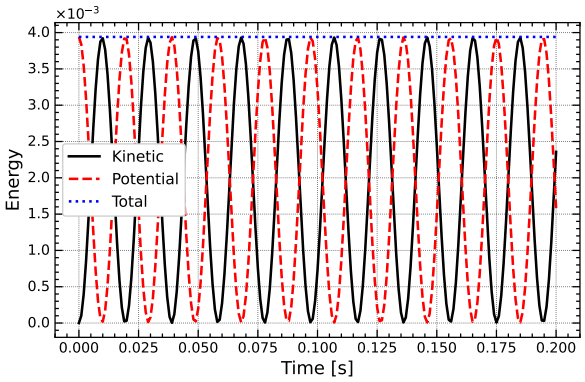
The response peak near the first natural frequency provides a complementary frequency-domain check of the assembled modal frequency scale and forcing.

**Localized patch-load response.** Under the upward-positive convention of Section 9.2, the verification driver applies a positive step load

$$q = 1500\,\mathrm{N\,m^{-2}}$$

over a $0.35\,\mathrm{m} \times 0.35\,\mathrm{m}$ patch centered at $(1.0, 0.5)\,\mathrm{m}$. This sign differs from the downward ROM-campaign loads, and the verification-driver

**(a) Undamped energy history.**




**(b) First-mode harmonic resonance.**

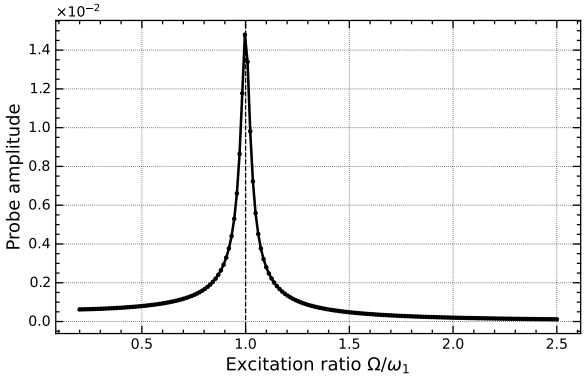



**Figure 10.3:** Energy- and frequency-domain verification at $N_\Omega = 32$ and $p = 2$. Panel (a) shows kinetic, potential, and total energy for foundation-free, undamped, simply supported free vibration. Panel (b) shows the first-mode harmonic response for the compression coefficient with the diagnostic modal damping ratio $\zeta = 0.02$.

patch is distinct from the $0.351\,\mathrm{m} \times 0.351\,\mathrm{m}$ patch stored in the campaign metadata. The prescribed foundation regime is compression.

The direct finite element history is compared with a twelve-mode finite element reconstruction. At the center probe, the peak absolute displacements are

$$\max_n \left| w^{\mathrm{direct}}_{P_c,h} \right| = 4.774513 \times 10^{-4}\,\mathrm{m}, \qquad \max_n \left| w^{\mathrm{modal}}_{P_c,h} \right| = 4.740118 \times 10^{-4}\,\mathrm{m}. \tag{10.75}$$

The history discrepancies are

$$e_{\ell^2} = 1.370831 \times 10^{-2}, \qquad e_{\ell^\infty} = 1.462636 \times 10^{-2}. \tag{10.76}$$

The direct and truncated-modal calls were also timed in the verification workflow. Because the absolute wall times vary across executions and hardware states, they are not used as thesis-level performance evidence; this test is interpreted through the field-history discrepancies above.

**Panel-partitioned and monolithic consistency.** A two-panel realization with one vertical interface is compared against the monolithic plate under the same simply supported boundary conditions, compression foundation coefficient, localized patch load, time grid, and material data. Before dynamic filtering, the panel-wise mixed system contains 9,542 algebraically free degrees of freedom. The mass-diagonal criterion in Equation (10.17) removes 4,718 zero-mass component degrees of freedom and leaves 4,824 active dynamic unknowns.

At the off-center probe $P_p$, the maximum absolute displacements are

$$\max_n \left| w^{\mathrm{mono}}_{P_p,h} \right| = 3.605595 \times 10^{-4}\,\mathrm{m}, \qquad \max_n \left| w^{\mathrm{panel}}_{P_p,h} \right| = 3.623223 \times 10^{-4}\,\mathrm{m}. \tag{10.77}$$

The relative history differences are

$$e_{\ell^2} = 4.063590 \times 10^{-3}, \qquad e_{\ell^\infty} = 5.785159 \times 10^{-3}. \tag{10.78}$$

This comparison simultaneously exercises the orientation-independent physical trace selection, restriction of $C^0$-IPG terms to same-panel facets, physical seam penalties, dynamic active-degree-of-freedom filtering, and lifting of the active panel state. The nonzero discrepancy is evaluated at finite interface weights and one fixed mesh.

**(a) Localized patch-load response.** **(b) Panel-to-monolithic consistency.**

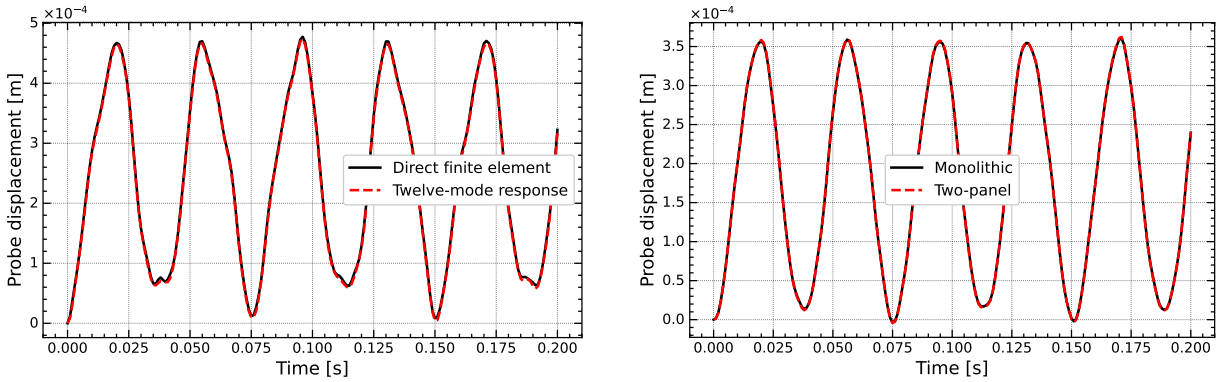



**Figure 10.4:** Verification under the positive localized step load $q = 1500\,\mathrm{N\,m^{-2}}$, with simply supported exterior edges, the prescribed compression coefficient, $N_\Omega = 32$, $p = 2$, and $\Delta\tau = 10^{-3}$ s. Panel (a) compares direct and twelve-mode monolithic histories at the center probe. Panel (b) compares monolithic and penalty-coupled two-panel histories at the off-center probe.

## Scope of the verification evidence

These checks establish internal consistency of the implemented operators and algorithms at the tested spatial and temporal resolutions [35, 36]. In particular, they check the frequency scale of the assembled stiffness and mass matrices in the isotropic specialization, the direct and modal Newmark implementations, the energy behavior of the undamped average-acceleration scheme, the sampled response near the first resonance, the localized forcing path, and the principal algebraic operations required by panel partitions.

These fixed-discretization checks are not experimental validation or asymptotic studies of spatial resolution, time step, or interface penalty, and they do not cover arbitrary anisotropic rigidity combinations. Within this scope, the results support the trajectories used in Chapter 11.

**Table 10.4:** Summary of the full-order verification program.

| Verification | Primary component tested | Reported metric |
|---|---|---|
| FEM versus Navier frequencies | Fourth-order bending stiffness, mass, effective foundation, and SSSS treatment | $\max e_\omega = 7.346886 \times 10^{-3}$ |
| Direct versus Navier single-mode Newmark | Load normalization, mass scaling, and direct time integration | $\max e_{\ell^2} = 1.314488 \times 10^{-2}$ |
| Direct versus eight-mode Newmark | Modal projection, load reduction, and reconstruction | $\max e_{\ell^2} = 8.470401 \times 10^{-7}$ |
| Undamped free vibration | Average-acceleration energy behavior | drift $1.039256 \times 10^{-10}$ |
| Harmonic sweep | First resonance location | relative error $3.351955 \times 10^{-3}$ |
| Localized patch load | Localized forcing and twelve-mode truncation | $e_{\ell^2} = 1.370831 \times 10^{-2}$ |
| Panel versus monolithic | Seam terms, physical traces, filtering, and lifting | $e_{\ell^2} = 4.063590 \times 10^{-3}$ |

## 10.7 Summary and handoff to the reduced-order model

The transient full-order approximation combines panel-wise quadratic Lagrange spaces, a symmetric $C^0$ interior-penalty bending form, separately assembled physical-seam penalties, a consistent transverse mass matrix, optional Rayleigh damping, and the average-acceleration Newmark scheme. The discrete construction preserves the distinction between numerical interelement stabilization and mechanical transmission across a physical seam. For mixed panel systems, the mass-diagonal filter removes zero-mass component entries and defines the active dynamic state used consistently by every subsequent algebraic operation.

At the fixed verification discretization, the maximum relative discrepancy between the first eight finite element and ordered Navier frequencies is $7.346886 \times 10^{-3}$. The direct and spatially analytical single-mode Newmark histories differ by at most $1.314488 \times 10^{-2}$ in relative discrete $\ell^2$ norm, whereas direct and eight-mode finite element Newmark solutions under deliberately single-mode forcing agree within $8.470401 \times 10^{-7}$. The undamped energy drift is $1.039256 \times 10^{-10}$; the sampled harmonic-response peak lies within $3.351955 \times 10^{-3}$ of the first natural frequency; the localized patch-load direct/modal discrepancy is $1.370831 \times 10^{-2}$; and the finite-penalty two-panel/monolithic probe-history difference is $4.063590 \times 10^{-3}$. Together, these results establish internal consistency at the tested settings.

The verified trajectories are high-dimensional and must be advanced over all time steps for every new physical query. Chapter 11 therefore transfers monolithic and panel-partitioned histories to a common continuous displacement output, compresses the stored fields by POD, and approximates the parameter–time-to-coordinate map with non-intrusive interpolation and regression models.

# Non-Intrusive Reduced-Order Models for Parameterized Transient Dynamics

# 11


11.1 Many-query objective and assessment structure . . 127
11.2 Time-dependent displacement output . . . . . 127
11.3 Snapshot campaigns and complete-trajectory data split . . . . . 128
11.4 Training-only POD construction and projection diagnostic . . . . . 130
11.5 Active non-intrusive coordinate models . . . . 131
11.6 Data partition, error measures, and timing protocol . . . . . 133
11.7 Reduced-order results . 134
11.8 Cross-campaign interpretation and conclusions . 140


## 11.1 Many-query objective and assessment structure

The verified full-order model of Chapter 10 remains expensive in many-query settings because every physical parameter value requires assembly of the active dynamic system and advancement through the complete Newmark time grid. This chapter therefore constructs non-intrusive reduced-order models for the transverse-displacement output of the transient problem. Spatial compression and parameter–time prediction are treated as distinct approximation tasks: Proper Orthogonal Decomposition (POD) supplies a low-dimensional displacement space, while interpolation, Gaussian-process, and neural-network surrogates approximate the map from the physical trajectory parameters and observation time to the retained POD coordinates [39–41, 53]. Time-dependent reduced-basis and non-intrusive ROM literature provides the broader context for separating spatial compression from learned parameter–time prediction [102–104].

The numerical protocol is trajectory-wise. For each campaign, complete trajectories are assigned once to disjoint training, validation, and test subsets before the stored time states are flattened. Training trajectories alone determine the POD basis, preprocessing transformations, interpolation and regression fits, and all gradient-based parameter updates. Held-out validation trajectories support early stopping, checkpoint selection, model-development decisions, and the aggregate cross-campaign assessment reported in this chapter. The test trajectories remain disjoint from all training, validation, preprocessing, model-fitting, and checkpoint-selection operations and do not enter the aggregate cross-campaign statistics. After all data-dependent choices have been fixed, one complete test trajectory per campaign is used only for a post-selection probe-history diagnostic, with the trajectory selected independently of ROM error. This construction prevents temporally adjacent states from one trajectory from crossing statistical subsets [54, 55].

The predictive models are PODI–RBF, PODI–Linear, POD–GPR, POD–NN, and POD–DL-ROM. POD projection is retained only as a compression reference because it requires the corresponding full-order field.

The numerical results report training-basis compression, retained-dimension sensitivity, held-out predictive field errors, online timing ratios, spatial error fields, and complete-trajectory probe histories. The timing quantity is a complete-FOM-trajectory-to-single-reduced-state cost ratio; the reserved test-trajectory probe histories are post-selection diagnostics rather than inputs to model ranking.

## 11.2 Time-dependent displacement output

Let $X_h \subset H^1(\Omega)$ denote the global continuous finite element space used for snapshot storage, POD construction, reconstruction, and physical

field-error evaluation. The high-fidelity displacement-output map is

$$\mathcal{S}_h : \mathcal{D} \times [0,T] \longrightarrow X_h, \qquad (\boldsymbol{\mu}, \tau) \longmapsto \widehat{w}_h(\boldsymbol{\mu}, \tau). \tag{11.1}$$

For a monolithic plate, $\widehat{w}_h$ is the finite element displacement represented directly in $X_h$. For a panel partition, the transient problem is solved in the mixed panel space introduced in Section 10.3. The active dynamic vector is first lifted to the complete mixed vector and reconstructed as a panel-wise field. If $\chi_i$ denotes the piecewise-constant indicator of $\Omega_i$, the common continuous output is

$$\widehat{w}_h = \mathcal{P}_h^{\mathrm{CG}} \left( \sum_{i=1}^{M} \chi_i w_{h,i} \right), \qquad \mathcal{P}_h^{\mathrm{CG}} : L^2(\Omega) \longrightarrow X_h, \tag{11.2}$$

where $\mathcal{P}_h^{\mathrm{CG}}$ is the discrete $L^2(\Omega)$ projection used to define the common output field. The transfer does not remove the panel-interface mechanics from the full-order solve. It provides a consistent continuous displacement field for the monolithic and panel-partitioned ROM campaigns. The displacement and normal-rotation jumps influence the panel-wise trajectories but are not retained as separate reduced outputs.

A physical trajectory is indexed by a parameter vector $\boldsymbol{\mu}^{(r)}$, $r = 1, \ldots, N_\mu$, and contains $N_\tau$ stored time states. The unflattened snapshot array is

$$\mathsf{W} \in \mathbb{R}^{N_\mu \times N_\tau \times N_h}, \qquad \mathsf{W}_{r,n,:} = \mathrm{dof}\Big( \widehat{w}_h(\boldsymbol{\mu}^{(r)}, \tau_n) \Big), \qquad N_h := \dim X_h. \tag{11.3}$$

For matrix compression and regression, the trajectory and time axes are flattened:

$$\begin{aligned} &\mathbf{W} \in \mathbb{R}^{N_{\mathrm{snap}} \times N_h}, \quad \mathbf{X} = \big[ (\boldsymbol{\mu}^{(r)}, \tau_n) \big]_{r,n} \in \mathbb{R}^{N_{\mathrm{snap}} \times (P+1)}, \\ &N_{\mathrm{snap}} = N_\mu N_\tau. \end{aligned} \tag{11.4}$$

Trajectory identifiers and stored-time indices are retained alongside the flattened rows. A physical trajectory, one stored state, one flattened snapshot row, an active dynamic degree of freedom, a mixed-space degree of freedom, and an output-space degree of freedom therefore remain distinct objects throughout the workflow.

## 11.3 Snapshot campaigns and complete-trajectory data split

All three campaigns use the full-order settings summarized in Table 11.1. The stored time vector is

$$\{\tau_n\}_{n=0}^{40} = \{0, 0.005, 0.010, \ldots, 0.200\} \text{ s}. \tag{11.5}$$

The plate has $L = 2\,\mathrm{m}$, $B = 1\,\mathrm{m}$, and $t = 0.05\,\mathrm{m}$. Each trajectory uses the downward rectangular patch load of Section 9.4, centered at $(1.0, 0.5)\,\mathrm{m}$ and having side lengths $0.351\,\mathrm{m}$. The exterior edges are free, the prescribed effective foundation regime is compression, the Rayleigh coefficients vanish, and the initial displacement and velocity are zero. The sampled load coordinate $A_q$ is positive, whereas the full-order parameter map applies the signed amplitude $-A_q$ under the upward-positive convention.

**Table 11.1:** Common full-order settings for the transient ROM campaigns. Time is an independent evolution variable and an input coordinate of the non-intrusive surrogates, not a material, geometric, or loading parameter.

| Quantity | Value used |
|---|---|
| Plate dimensions | $L = 2\,\mathrm{m}$, $B = 1\,\mathrm{m}$, $t = 0.05\,\mathrm{m}$ |
| Resolution selector / polynomial degree | $N_\Omega = 64$ / $p = 2$ |
| Final time / integration step | $T = 0.20\,\mathrm{s}$ / $\Delta\tau = 10^{-3}\,\mathrm{s}$ |
| Newmark parameters | $\beta = 1/4$, $\gamma = 1/2$ |
| Rayleigh coefficients | $\alpha_M = 0\,\mathrm{s}^{-1}$, $\alpha_K = 0\,\mathrm{s}$ |
| Integration steps / storage interval | 200 / every five steps |
| Stored states per trajectory | 41 |
| Load | downward rectangular patch with step temporal profile |
| Patch center / dimensions | $(1.0, 0.5)\,\mathrm{m}$ / $0.351\,\mathrm{m} \times 0.351\,\mathrm{m}$ |
| Exterior boundary / foundation regime | free / prescribed compression coefficient |
| Initial data | $w_0 = 0$, $v_0 = 0$ |

**Table 11.2:** Trajectory databases, disjoint complete-trajectory partitions, continuous-output dimensions, and final POD dimensions. The row counts are obtained by multiplying each trajectory count by the 41 stored times.

| Configuration | Input | $N_\mu$ | $N_\tau$ | Trajectory split | Row split | $\dim X_h / N$ |
|---|---|---|---|---|---|---|
| Case 2 monolithic | $(A_q, \tau)$ | 200 | 41 | 160/20/20 | 6560/820/820 | 20,503/7 |
| Case 2 contact 1–0 | $(A_q, \tau)$ | 250 | 41 | 200/25/25 | 8200/1025/1025 | 20,375/8 |
| Case 3 monolithic | $(D_x, A_q, \tau)$ | 250 | 41 | 200/25/25 | 8200/1025/1025 | 20,503/8 |

The three databases and their complete-trajectory splits are listed in Table 11.2. Case 2 monolithic varies only the load amplitude and uses the extended input $\xi = (A_q, \tau)$. Case 2 contact 1–0 uses the same reduced input but solves a two-panel configuration with one vertical physical interface. Case 3 monolithic varies the linked rigidity/load coordinates and uses $\xi = (D_x, A_q, \tau)$. In the last campaign, $D_x$ is the sampled rigidity coordinate; the remaining orthotropic rigidities are reconstructed through the same deterministic parameterization used during snapshot generation and are not independent inputs.

The parameter intervals below are the extrema present in the stored campaigns. For Case 2 monolithic,

$$A_q \in [2008.7, 9972.1]\,\mathrm{N\,m^{-2}}. \tag{11.6}$$

For Case 3 monolithic,

$$D_x \in [4015.1, 13982]\,\mathrm{N\,m}, \qquad A_q \in [2010.5, 9986.4]\,\mathrm{N\,m^{-2}}. \tag{11.7}$$

For Case 2 contact 1–0,

$$A_q \in [2006.3, 9981.3]\,\mathrm{N\,m^{-2}}. \tag{11.8}$$

The two-panel full-order realization contains 40,750 algebraic mixed-space degrees of freedom before dynamic filtering, 20,488 active dynamic unknowns after removal of the zero-mass entries, and 20,375 degrees of freedom in the continuous output space used by the ROM.

The split is applied to unique trajectory identifiers rather than to flattened space–time rows. Seed 100 selects 80% of the trajectories for training, and seed 101 divides the remainder equally between validation and testing. If $\mathcal{I}^{\mu}_{\mathrm{tr}}$, $\mathcal{I}^{\mu}_{\mathrm{val}}$, and $\mathcal{I}^{\mu}_{\mathrm{te}}$ denote the trajectory-index sets, then

$$\mathcal{I}^{\mu}_{\mathrm{tr}} \,\dot\cup\, \mathcal{I}^{\mu}_{\mathrm{val}} \,\dot\cup\, \mathcal{I}^{\mu}_{\mathrm{te}} = \{1, \dots, N_{\mu}\}, \qquad \mathcal{I}^{\mu}_{a} \cap \mathcal{I}^{\mu}_{b} = \varnothing \quad (a \neq b). \tag{11.9}$$

Every stored state associated with one trajectory inherits the same assignment. Training trajectories determine POD construction and predictive-model fitting. Held-out validation trajectories support model development and the reported aggregate cross-campaign assessment. Test trajectories remain reserved from model development and do not enter the aggregate accuracy statistics or method ranking.

## 11.4 Training-only POD construction and projection diagnostic

Let $\mathbf{W}_{\mathrm{tr}}$ denote the uncentered snapshot matrix formed only from rows belonging to the training trajectories. Its leading Euclidean singular triplets satisfy

$$\mathbf{W}_{\mathrm{tr}} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^{\top}, \qquad \boldsymbol{\Sigma} = \mathrm{diag}(\sigma_1, \dots, \sigma_r), \qquad \sigma_1 \geq \sigma_2 \geq \cdots \geq 0. \tag{11.10}$$

Thirty singular modes are retained for the spectral analysis. No mean field is subtracted, so the spectral extraction is an uncentered Euclidean compression of the training snapshot matrix. Snapshot-based POD is widely used for reduction of time-dependent PDEs and for dynamical characterization of mechanical systems [105, 106].

The spatial vectors obtained from $\mathbf{V}$ are converted to finite element functions and reorthonormalized in the discrete $H^1$-type product

$$(u, v)_X := \int_{\Omega} uv \,\mathrm{d}\Omega + \int_{\Omega} \nabla u \cdot \nabla v \,\mathrm{d}\Omega. \tag{11.11}$$

Let $\mathbf{G}_X$ be the associated Gram matrix and let $\boldsymbol{\Phi}_N \in \mathbb{R}^{N_h \times N}$ contain the reorthonormalized mode vectors. Then

$$\boldsymbol{\Phi}_N^{\top} \mathbf{G}_X \boldsymbol{\Phi}_N = \mathbf{I}_N \tag{11.12}$$

up to the numerical orthonormalization tolerance. More generally, with

$$\mathbf{G}_N := \boldsymbol{\Phi}_N^{\top} \mathbf{G}_X \boldsymbol{\Phi}_N,$$

the POD coordinates of an available displacement vector $\widehat{\mathbf{w}}_h$ are defined by

$$\mathbf{G}_N \mathbf{a}_N = \boldsymbol{\Phi}_N^{\top} \mathbf{G}_X \widehat{\mathbf{w}}_h. \tag{11.13}$$

The resulting reduced space and projector are

$$X_N = \mathrm{span}\{\phi_1, \dots, \phi_N\}, \qquad \mathcal{P}_N^X \widehat{\mathbf{w}}_h = \boldsymbol{\Phi}_N \mathbf{G}_N^{-1} \boldsymbol{\Phi}_N^{\top} \mathbf{G}_X \widehat{\mathbf{w}}_h. \tag{11.14}$$

The implemented construction therefore combines Euclidean spectral extraction with finite element reorthonormalization and projection in the product (11.11); it is not a purely $H^1$-optimal POD.

For an available full-order state with $\|\widehat{w}_h(\boldsymbol{\mu}, \tau)\|_{L^2(\Omega)} > 0$, the training-basis POD projection benchmark is evaluated in the physical finite

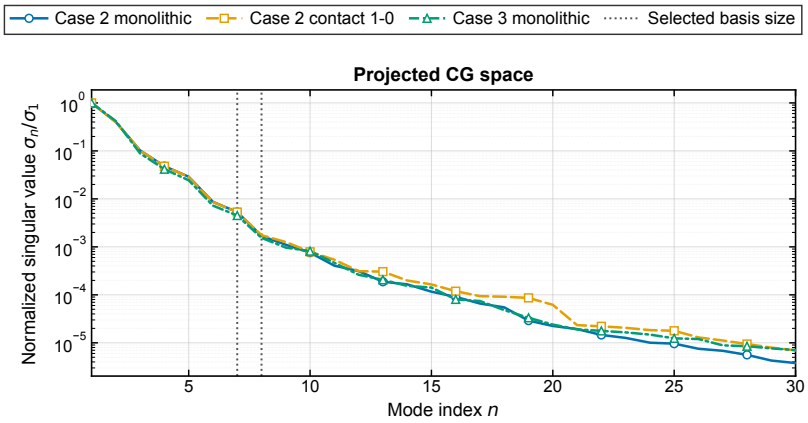



**Figure 11.1:** Training-only normalized singular-value decay in the projected continuous displacement space. Vertical lines mark $N = 7$ for Case 2 monolithic and $N = 8$ for the two-panel Case 2 and Case 3 monolithic campaigns.

element $L^2(\Omega)$ norm,

$$e_{\mathrm{proj}}(\boldsymbol{\mu}, \tau; N) := \frac{\left\| \widehat{w}_h(\boldsymbol{\mu}, \tau) - \mathcal{P}_N^X \widehat{w}_h(\boldsymbol{\mu}, \tau) \right\|_{L^2(\Omega)}}{\left\| \widehat{w}_h(\boldsymbol{\mu}, \tau) \right\|_{L^2(\Omega)}}. \tag{11.15}$$

In the reported campaigns, the zero-time state satisfies $\widehat{w}_h(\boldsymbol{\mu}, 0) = 0$; its relative projection error is therefore undefined and is excluded from relative-error aggregation together with the corresponding predictive relative errors. The quantity in Equation (11.15) is a compression diagnostic rather than an online predictor because it requires the full-order field. Moreover, the projection is defined in the $H^1$-type product of Equation (11.11), whereas the assessment uses the $L^2(\Omega)$ norm. It is therefore a reference benchmark rather than a strict lower bound on the reported predictive error.

The normalized singular values and cumulative uncentered snapshot energy are

$$\widehat{\sigma}_n = \frac{\sigma_n}{\sigma_1}, \qquad n = 1, \ldots, r, \tag{11.16}$$

and

$$\mathcal{E}_N = \frac{\sum_{j=1}^{N} \sigma_j^2}{\sum_{j=1}^{r} \sigma_j^2}. \tag{11.17}$$

The three spectra have nearly identical rapid initial decay. In every campaign, the first four modes contain more than 99.9% of the uncentered Euclidean training-snapshot energy. The final dimensions $N = 7, 8, 8$ therefore exceed the corresponding 99.9% thresholds and provide additional resolution beyond the dominant spectral range. Figure 11.1 establishes compression of the training displacement manifolds; it does not establish the accuracy of the parameter–time coordinate maps.

## 11.5 Active non-intrusive coordinate models

For a fixed basis dimension $N$, each predictive method approximates

$$\mathcal{R}_N : \boldsymbol{\xi} = (\boldsymbol{\mu}, \tau) \longmapsto \widetilde{\mathbf{a}}_N(\boldsymbol{\xi}) \in \mathbb{R}^N, \qquad \widetilde{w}_N(\boldsymbol{\xi}) = \sum_{j=1}^{N} \widetilde{a}_j(\boldsymbol{\xi}) \phi_j. \tag{11.18}$$

| Method | Learned map | Principal setting |
|---|---|---|
| PODI–RBF | $\xi \mapsto \widetilde{\mathbf{a}}_N$ | local thin-plate-spline RBF; smoothing $10^{-10}$; 96 neighbors |
| PODI–Linear | $\xi \mapsto \widetilde{\mathbf{a}}_N$ | Delaunay interpolation; at most 2500 fit states; nearest-state fallback |
| POD–GPR | $\xi \mapsto \widetilde{a}_j$, independently | standardized Matérn-5/2 GP plus white noise; at most 1200 states |
| POD–NN | $\xi \mapsto \widetilde{\mathbf{a}}_N$ | ELU network $(128, 96, 64)$; coefficient MSE; validation checkpointing |
| POD–DL-ROM | enriched $\xi \mapsto \widetilde{\mathbf{a}}_N$ | SiLU network $(192, 160, 128)$; train-only enriched features, layer normalization, dropout 0.02 |

**Table 11.3:** Active predictive transient ROMs and their principal settings.

Input transformations and, where used, coordinate standardization are fitted from the training data only. The known initial condition is imposed at prediction time,

$$\widetilde{\mathbf{a}}_N(\boldsymbol{\mu}, 0) = \mathbf{0}. \tag{11.19}$$

The zero-time snapshots remain in the databases; Equation (11.19) enforces the physical initial state rather than deleting those rows.

## POD interpolation

PODI–RBF standardizes the active inputs and POD coordinates and applies a local thin-plate-spline radial-basis interpolant with smoothing $10^{-10}$, polynomial degree one, and 96 eligible neighbors per query [81, 82, 104]. PODI–Linear constructs a Delaunay piecewise-linear interpolant in the normalized parameter–time space. Its fit set contains at most 2500 states to control tessellation and memory costs. When trajectory identifiers are available, the retained fit states are selected through trajectory-aware coverage; otherwise, a state-wise fallback is used. Queries outside the resulting convex hull return the coordinate vector of the nearest fit state.

## Gaussian-process regression

POD–GPR fits one independent scalar Gaussian process to each standardized POD coordinate [77, 103]. The reported held-out validation assessment uses a Matérn-5/2 covariance kernel augmented by a white-noise contribution, diagonal regularization $10^{-9}$, and two optimizer restarts. At most 1200 training states are selected through a complete-time, maximin strategy to limit cubic fitting cost. The posterior mean supplies the predicted coordinates. Posterior variances are available but are not propagated to probabilistic displacement intervals in this chapter.

## POD neural-network regression

POD–NN is a fully connected coordinate regressor with hidden widths $(128, 96, 64)$, ELU activations, and a linear output [47, 107]. Inputs and POD projection coordinates are standardized from the training trajectories. The network is trained directly on the precomputed standardized coordinates by minimizing the coefficient mean-square error with

AdamW optimization, validation-based learning-rate reduction, early stopping, and restoration of the minimum-validation-loss checkpoint. At prediction time, the coordinates are inverse-scaled, the prescribed initial state is imposed through Equation (11.19), and the displacement is reconstructed through Equation (11.18).

### POD-enhanced deep-learning ROM

POD–DL-ROM enriches the parameter–time input before predicting the POD coordinates, following the general principle of POD-assisted nonlinear reduction [49]. The train-fitted feature map concatenates the raw query coordinates, their min–max-normalized copies, clipped log-normalized copies, the first four powers of normalized time, and six sine–cosine harmonic pairs. The resulting feature vector has dimension 22 for the two Case 2 campaigns and 25 for Case 3 monolithic and is subsequently standardized from the training subset.

A sequential (192, 160, 128) SiLU network with layer normalization and dropout 0.02 maps these features to the retained displacement POD coordinates. The feature-to-coordinate network is the deployed online predictor, while the coefficient autoencoder supports auxiliary training diagnostics. Training includes reconstruction-level refinement after POD lifting, and the prescribed zero-time coordinate vector is imposed before field reconstruction. The resulting model is a direct time-conditioned coordinate surrogate. The complete executed architecture and principal hyperparameters are documented in Section A.5.

## 11.6 Data partition, error measures, and timing protocol

The trajectory partition defined in Equation (11.9) is used unchanged throughout the assessment. The training-only operations and the roles assigned to the validation and test trajectories are those stated in Section 11.3. In particular, test trajectories remain excluded from model development and enter only the post-selection probe-history diagnostic described below.

For a held-out validation state $j$, let $\widehat{\mathbf{w}}_h^{(j)}$ and $\widetilde{\mathbf{w}}_N^{(j)}$ denote the full-order and reconstructed displacement coefficient vectors in $X_h$, and let $\mathbf{M}_h$ be the finite element mass matrix associated with the common output space. The relative physical field error is

$$e_{L^2}^{(j)} := \frac{\sqrt{\left(\widetilde{\mathbf{w}}_N^{(j)} - \widehat{\mathbf{w}}_h^{(j)}\right)^\top \mathbf{M}_h \left(\widetilde{\mathbf{w}}_N^{(j)} - \widehat{\mathbf{w}}_h^{(j)}\right)}}{\sqrt{\left(\widehat{\mathbf{w}}_h^{(j)}\right)^\top \mathbf{M}_h \widehat{\mathbf{w}}_h^{(j)}}}. \tag{11.20}$$

The associated absolute field norm is

$$e_{L^2,\mathrm{abs}}^{(j)} := \sqrt{\left(\widetilde{\mathbf{w}}_N^{(j)} - \widehat{\mathbf{w}}_h^{(j)}\right)^\top \mathbf{M}_h \left(\widetilde{\mathbf{w}}_N^{(j)} - \widehat{\mathbf{w}}_h^{(j)}\right)}. \tag{11.21}$$

All trajectories start from the exact zero displacement, so the raw relative error at $\tau = 0$ is mathematically undefined. Those entries are retained as missing values and excluded from the mean, median, and maximum

relative-error statistics. All nonzero-time reference norms are finite and positive.

The timing assessment records

$$R^{(j)}_{\text{traj/state}} := \frac{t^{(j)}_{\text{FOM,traj}}}{t^{(j)}_{\text{ROM,state}}}. \tag{11.22}$$

The numerator is the wall time of a complete full-order trajectory. The denominator is the time required to transform one parameter–time query, predict one coordinate vector, and reconstruct one finite element displacement field. It excludes snapshot generation, POD construction, model fitting, checkpoint loading, file input/output, and plotting. Since the numerator and denominator correspond to different output requests, $R_{\text{traj/state}}$ is an online cost ratio rather than a matched end-to-end speedup. A matched trajectory-level speedup would instead require

$$S_{\text{traj}} = \frac{t_{\text{FOM,traj}}}{t_{\text{ROM,traj}}}, \qquad t_{\text{ROM,traj}} = \text{time}\left(\{\widetilde{w}_N(\boldsymbol{\mu}, \tau_n)\}_{n=0}^{N_\tau - 1}\right), \tag{11.23}$$

which has not been evaluated in this chapter. Accordingly, no speedup figure is reported for the transient ROM assessment.

## 11.7 Reduced-order results

### Reduced-dimension sensitivity

Figure 11.2 provides a complementary basis-enrichment study for the Case 2 monolithic load-amplitude map. The curve for each method is the mean relative $L^2(\Omega)$ displacement error over the diagnostic sample set, and the shaded band spans the corresponding sample-wise minimum and maximum. The figure is used only to illustrate dimension sensitivity. The primary cross-campaign evidence is the trajectory-held-out validation assessment at the fixed dimensions of Table 11.2.

The projection and interpolation curves decrease rapidly as the basis is enriched, showing that basis truncation dominates at the smallest dimensions. At larger dimensions, the coordinate surrogates reach method-dependent plateaus. POD–NN is the most visibly non-monotone, while POD–DL-ROM continues to improve but remains above the interpolation and GPR curves in this smooth load-only campaign. The result reinforces that basis enrichment and coordinate-map approximation must be assessed separately.

### Complete-trajectory held-out validation accuracy

The selected models are assessed on 20 held-out validation trajectories in Case 2 monolithic and 25 held-out validation trajectories in each of the other campaigns. Since 41 states are stored per trajectory, the corresponding validation subsets contain 820, 1025, and 1025 space–time rows. The 20, 25, and 25 zero-time rows are omitted only from aggregate relative-error statistics, as explained in Section 11.6.

The two load-amplitude maps are the least demanding. In Case 2 monolithic, PODI–Linear gives the lowest predictive mean, $3.421914 \times 10^{-3}$,

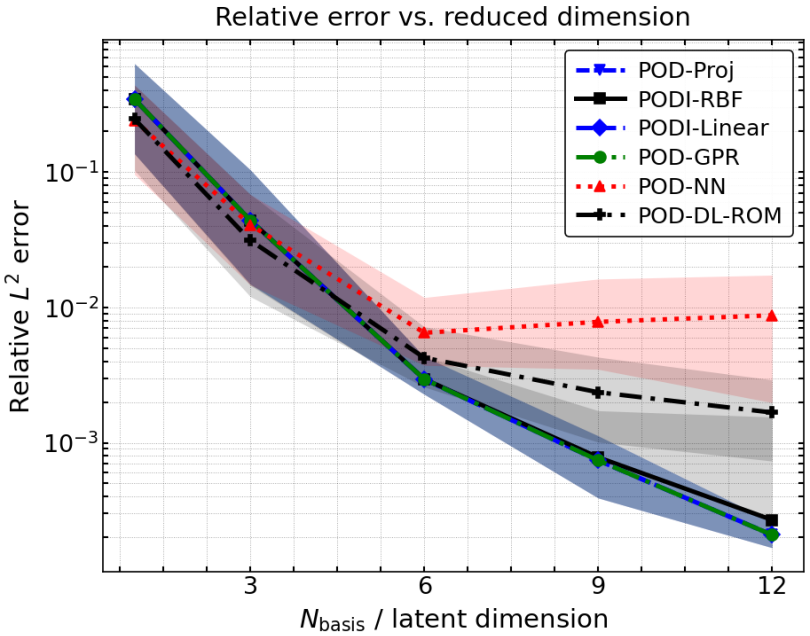



**Figure 11.2:** Case 2 monolithic reduced-dimension sensitivity: mean relative $L^2(\Omega)$ displacement error with sample-wise minimum–maximum bands. POD-Proj shows the corresponding training-basis projection reference.

while POD–GPR gives $3.421927 \times 10^{-3}$; both are effectively coincident with the training-basis projection benchmark at the reported precision. In the two-panel Case 2 campaign, POD–GPR gives the lowest predictive mean, $2.463655 \times 10^{-3}$, again essentially equal to the projection benchmark. These results show that the panel-wise solve, active-degree-of-freedom filtering, continuous output transfer, training-only POD compression, and non-intrusive reconstruction can be combined without introducing a separately learned interface coordinate.

Case 3 monolithic is more demanding because variation of the linked rigidity entries changes the modal spectrum and the phase accumulated over the transient. The POD projection benchmark remains small, $1.615441 \times 10^{-3}$, but no predictive method reaches it. POD–DL-ROM is the most accurate predictive model, with mean, median, and maximum errors $1.6749 \times 10^{-2}$, $1.3464 \times 10^{-2}$, and $1.2626 \times 10^{-1}$, respectively. The gap between the projection benchmark and the predictive curves shows that coordinate prediction, rather than linear state compression, controls the

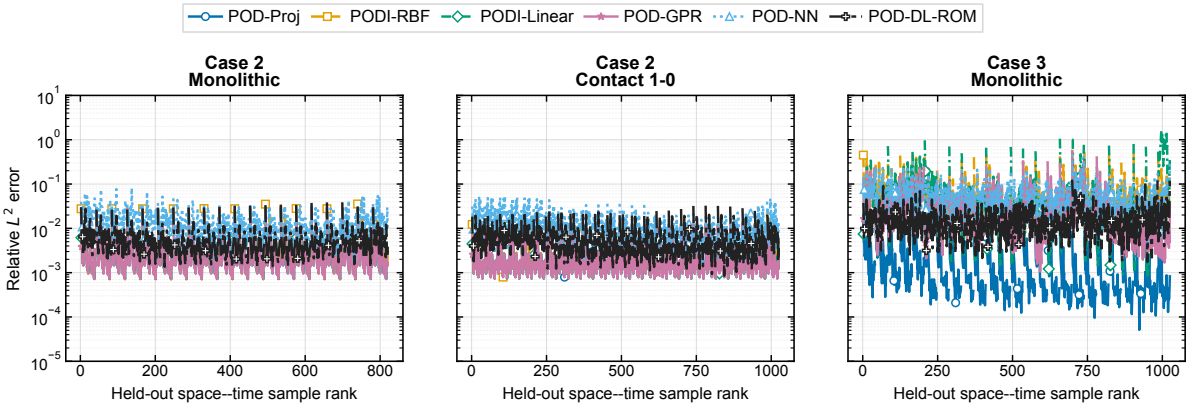



**Figure 11.3:** Held-out validation relative $L^2(\Omega)$ displacement errors for Case 2 monolithic, two-panel Case 2, and Case 3 monolithic. Nonzero-time states are ordered by held-out validation space–time sample rank. Predictive methods are shown together with the POD-Proj compression reference.

**Table 11.4:** Mean relative $L^2(\Omega)$ displacement errors over the nonzero-time held-out validation states. Bold entries denote the lowest predictive mean.

| Method | Case 2 monolithic | Case 2 contact 1–0 | Case 3 monolithic |
|---|---|---|---|
| POD-Proj | $3.421914 \times 10^{-3}$ | $2.463648 \times 10^{-3}$ | $1.615441 \times 10^{-3}$ |
| PODI–RBF | $3.772277 \times 10^{-3}$ | $2.623313 \times 10^{-3}$ | $7.016461 \times 10^{-2}$ |
| PODI–Linear | $\mathbf{3.421914 \times 10^{-3}}$ | $2.489550 \times 10^{-3}$ | $9.316652 \times 10^{-2}$ |
| POD–GPR | $3.421927 \times 10^{-3}$ | $\mathbf{2.463655 \times 10^{-3}}$ | $3.455144 \times 10^{-2}$ |
| POD–NN | $1.535103 \times 10^{-2}$ | $1.327236 \times 10^{-2}$ | $6.260257 \times 10^{-2}$ |
| POD–DL-ROM | $6.044031 \times 10^{-3}$ | $6.264467 \times 10^{-3}$ | $\mathbf{1.674858 \times 10^{-2}}$ |

error in this campaign. By contrast, the best predictive maxima in Case 2 monolithic and Case 2 contact 1–0 are $2.7717 \times 10^{-2}$ and $1.2084 \times 10^{-2}$, respectively.

## Online cost

The mean trajectory-to-single-state ratios of Equation (11.22) range from approximately $1.86 \times 10^4$ to $6.39 \times 10^4$ over the active predictive methods and three campaigns. The smallest mean occurs for POD–GPR in Case 3 monolithic, while the largest occurs for PODI–Linear in Case 2 contact 1–0. These values show that one reconstructed reduced state is inexpensive relative to one complete full-order trajectory. They are not interpreted as matched trajectory-level speedups, and no speedup plot is included.

## Spatial reconstruction and absolute-error fields

The sample-wise curves are complemented by spatial diagnostics evaluated at one selected held-out validation space–time state per campaign. Within each campaign, the same state is used for POD–Proj and all five predictive ROMs. Among the available candidate states, the displayed state is the one with the lowest mean error across the five predictive ROMs and is therefore explicitly a best-performing spatial illustration. Aggregate validation behavior is reported in Figure 11.3 and Table 11.4.

The projected solutions retain the same global displacement pattern across the three cases, whereas the absolute-error fields reveal different spatial truncation structures. The panel case is evaluated after transfer to the global continuous output space, so the plotted projection error has the same mathematical definition as in the monolithic cases, although the color scales differ. In the two load-amplitude campaigns, the interpolation and GPR error fields closely follow the POD truncation structure, consistent with their aggregate accuracy. POD–NN and POD–DL-ROM show larger but spatially organized residuals. In the rigidity/load campaign, the five methods exhibit visibly different error patterns and amplitudes. POD–DL-ROM has the smallest aggregate predictive error, while the larger GPR and interpolation fields reflect the greater difficulty of mapping rigidity-dependent frequency and phase changes to POD coordinates.

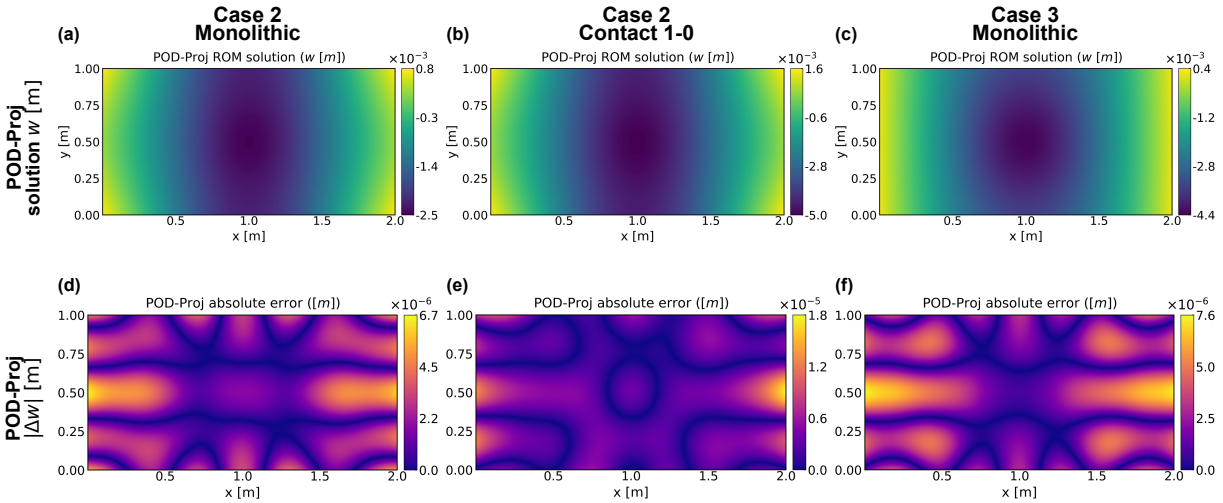



**Figure 11.4:** POD-Proj displacement (top) and absolute error (bottom) for the selected held-out validation state in each campaign. Columns show Case 2 monolithic, two-panel Case 2, and Case 3 monolithic; color scales are panel-specific. Aggregate validation accuracy is reported in Figure 11.3 and Table 11.4.

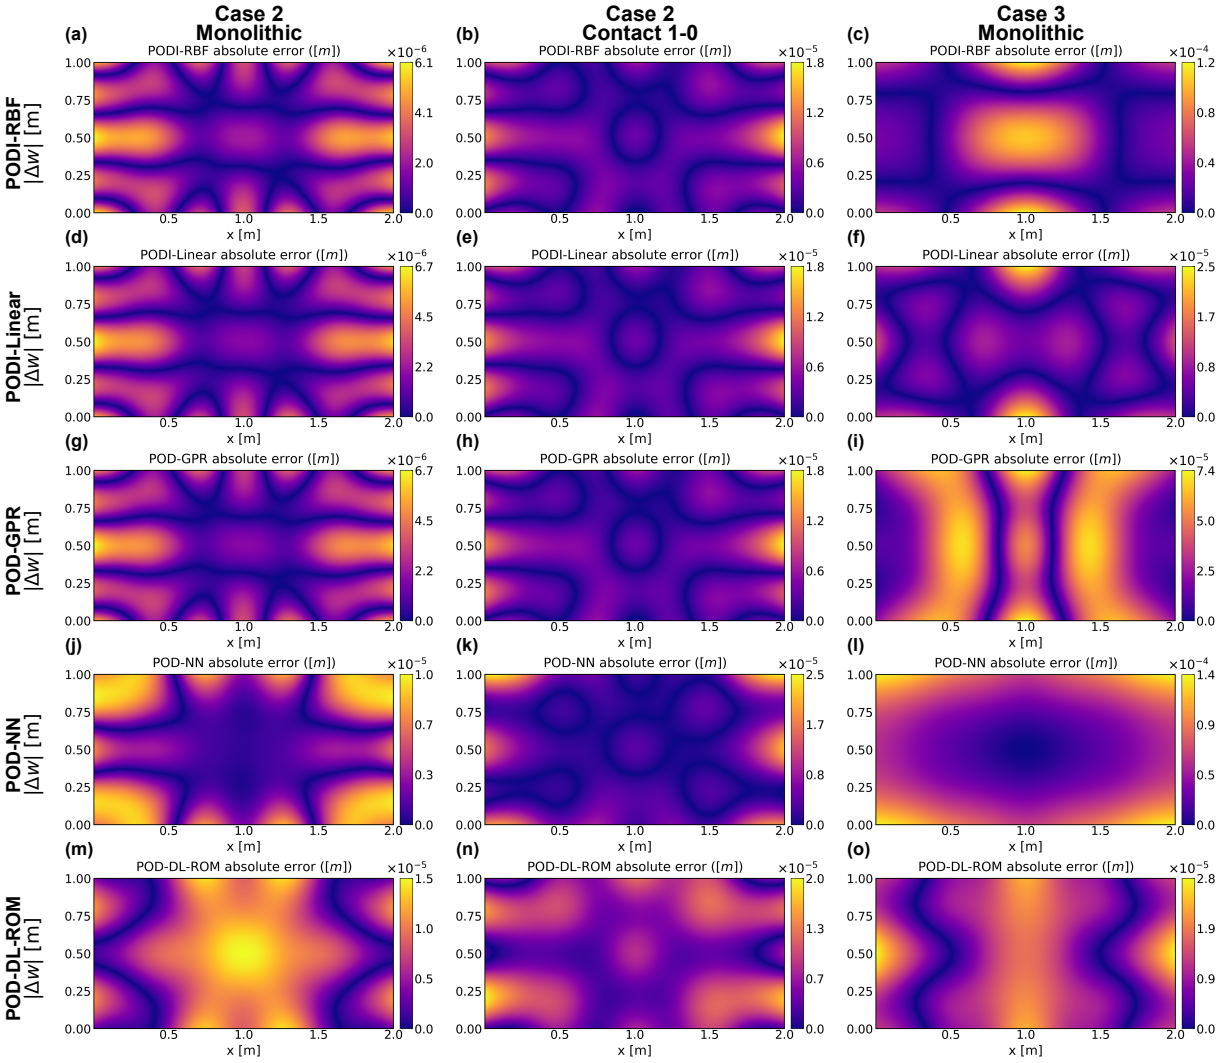



**Figure 11.5:** Predictive absolute-error fields at the selected held-out validation state in each campaign. Rows show PODI–RBF, PODI–Linear, POD–GPR, POD–NN, and POD–DL-ROM; columns show Case 2 monolithic, two-panel Case 2, and Case 3 monolithic. Color scales are panel-specific; aggregate validation accuracy is reported in Figure 11.3 and Table 11.4.

## Complete held-out test-trajectory probe histories and errors

The aggregate held-out validation statistics reported above quantify the finite element field error over all nonzero stored validation states, while the spatial diagnostics identify the structure of the displacement error at selected space–time states. Those measures do not directly show whether the predicted oscillations remain synchronized with the full-order response over one complete trajectory. Figure 11.6 therefore provides a complementary temporal diagnostic by comparing the FOM and ROM responses at the fixed probe $\boldsymbol{P} = (\mathbf{0.73}, \mathbf{0.37})\,\mathbf{m}$ for one complete held-out test trajectory from each campaign. The left column shows the probe-displacement histories, and the right column shows the corresponding absolute probe errors.

The displayed trajectories are evaluated only after the training-based POD spaces, preprocessing transformations, retained dimensions, and validation-selected model checkpoints have been fixed. Their selection is independent of the ROM errors displayed below. They therefore provide post-selection trajectory illustrations rather than additional model-selection evidence or aggregate test-set statistics. Each trajectory contains the same 41 stored states over $0 \leq \tau \leq 0.20\,\mathrm{s}$.

The two load-amplitude campaigns show close temporal agreement between the full-order and reduced probe responses. In both Case 2 histories, the principal displacement extrema, reversal times, and oscillatory phase are reproduced over the complete interval. At the displacement scale, most ROM curves are nearly superposed with the FOM, demonstrating why the accompanying absolute-error panels are required: small method-dependent discrepancies that are visually concealed in the displacement histories become distinguishable when plotted directly as errors.

For the monolithic Case 2 trajectory, the largest visible error excursions are localized in time rather than accumulated progressively over the trajectory. POD–Proj, PODI–Linear, POD–GPR, and the remaining predictive reconstructions reproduce the global waveform closely, while POD–NN shows the largest isolated probe-error peaks. This local behavior does not replace the field-wise validation statistics reported above, but it shows when the coordinate-prediction discrepancy becomes most pronounced during an otherwise accurately reconstructed trajectory.

The two-panel Case 2 result confirms that the same time-conditioned prediction strategy remains effective for the penalty-coupled configuration. The FOM and ROM histories retain close agreement over the complete response, although the error panel exposes larger method-dependent fluctuations for POD–NN and POD–DL-ROM than for POD–GPR. The comparison is performed in the transferred global continuous displacement space used throughout the ROM assessment and at a probe located away from the physical interface. It therefore assesses temporal prediction of the transferred displacement output, but it does not independently quantify displacement jumps, normal-rotation jumps, interface resultants, or other seam-local quantities.

The Case 3 monolithic trajectory exhibits a broader method-dependent spread in its absolute-error history. In particular, the paired representation makes visible an early PODI–Linear discrepancy and later POD–NN excursions that remain difficult to distinguish in the displacement panel itself. The principal oscillatory pattern is nevertheless retained by all

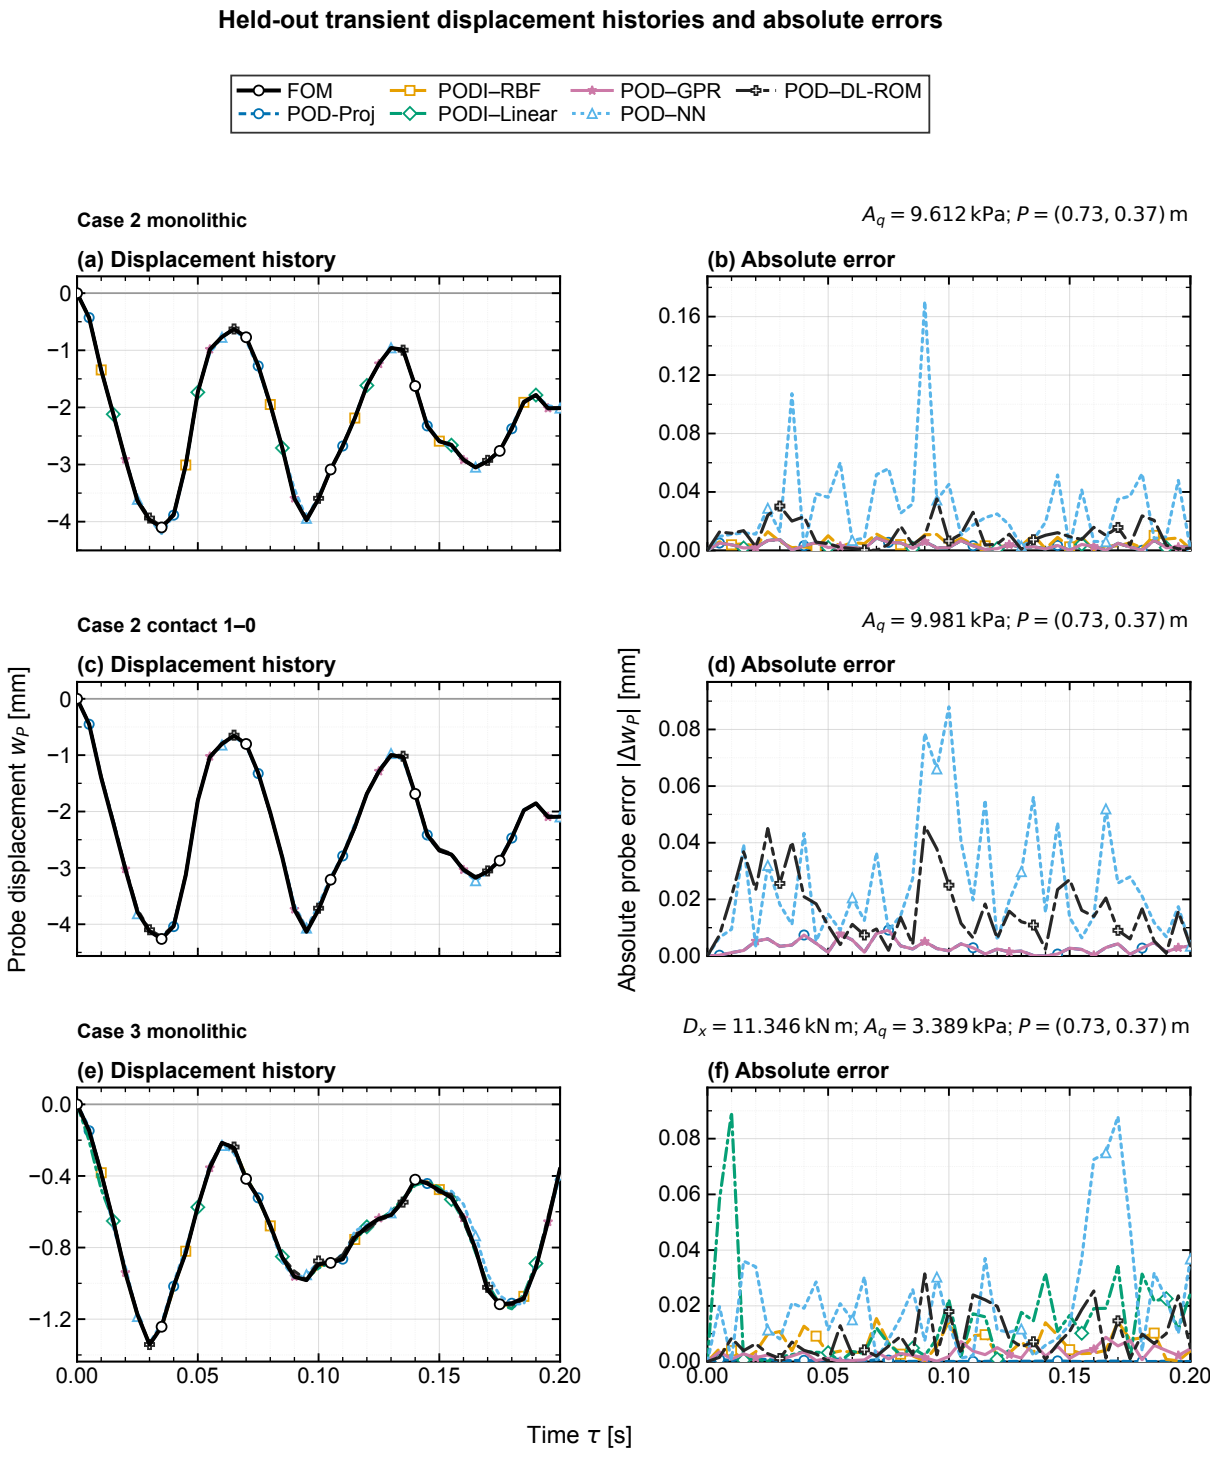



**Figure 11.6:** Complete held-out test-trajectory probe histories and absolute errors at $P = (0.73, 0.37)$ m. The left column shows the FOM and available ROM probe-displacement histories, and the right column shows $|w_P^{\mathrm{ROM}}(\tau) - w_P^{\mathrm{FOM}}(\tau)|$ in millimeters. Panels **(a,b)** correspond to Case 2 monolithic with $A_q = 9.612$ kPa; panels **(c,d)** to the two-panel Case 2 contact 1–0 configuration with $A_q = 9.981$ kPa; and panels **(e,f)** to Case 3 monolithic with $D_x = 11.346$ kN m and $A_q = 3.389$ kPa. POD–Proj is the non-predictive retained-space compression diagnostic, whereas the remaining curves are predictive ROM results. The shared legend reports the union of methods available across the three exports; PODI–RBF and PODI–Linear are unavailable for the two-panel Case 2 trajectory and therefore do not appear in panels **(c,d)**. All panels use the same stored time grid and time limits, while independent ordinate ranges preserve the visibility of the case-specific response amplitudes and smaller temporal errors.

displayed reconstructions. This behavior is consistent with the aggregate validation assessment: variation of $D_x$ changes the modal frequencies and

hence the accumulated phase, making the $(D_x, A_q, \tau)$-to-coordinate map more demanding than the load-amplitude-only maps. The aggregate finite element field-error statistics, rather than this single probe trajectory, remain the basis for identifying POD–DL-ROM as the most accurate predictive method for the rigidity-dependent campaign.

Figure 11.6 should consequently be read together with the relative $L^2(\Omega)$ statistics and the spatial error maps. It supplies a time-resolved local diagnostic at one observation point and demonstrates whether waveform, extrema, and phase remain aligned over a complete held-out trajectory.

## 11.8 Cross-campaign interpretation and conclusions

The three campaigns distinguish spatial compressibility from the difficulty of the parameter–time coordinate map. Their training singular values decay at similar rates, and four modes capture more than 99.9% of the uncentered Euclidean training-snapshot energy in every campaign. Predictive accuracy, however, depends on how the active parameters modify the trajectory. In the two load-amplitude studies, PODI–Linear and POD–GPR reproduce the held-out validation trajectories close to the projection benchmark because the dynamic operator is fixed. In the rigidity-dependent campaign, stiffness changes the modal spectrum and accumulated phase; POD–DL-ROM gives the lowest predictive mean error, $1.674858 \times 10^{-2}$, while the projection benchmark is $1.615441 \times 10^{-3}$. The remaining discrepancy is therefore associated primarily with the learned parameter–time map rather than with the retained spatial subspace.

All reported reductions preserve the trajectory-wise partition of Section 11.3. The retained dimensions are $N = 7$ for Case 2 monolithic and $N = 8$ for the two-panel Case 2 and Case 3 monolithic campaigns, with displacement as the reduced output and time as an independent query coordinate.

The reported evidence concerns the sampled parameter ranges, the 0.20 s horizon, step loading, free exterior edges, a prescribed compression-regime foundation coefficient, zero Rayleigh damping, and displacement prediction. The online ratio compares one complete full-order trajectory with one reduced-state query. Longer horizons, additional loading and damping families, adaptive sampling, independent trajectory campaigns, and matched complete-trajectory timing provide natural extensions of the present assessment.

# Conclusions and Future Outlook

This dissertation develops three parameterized orthotropic-plate studies covering static mechanics, steady thermomechanical bending, and transient response. They use the same Kirchhoff–Love kinematics, aligned orthotropic bending, compliant-support description, and panel-wise geometry, but retain distinct state variables, coupling mechanisms, and numerical assessments. In each study, the continuous model is linked to its high-fidelity approximation and then to the reduced output used for many-query prediction. A common conclusion emerges: low-dimensional state structure is not sufficient by itself to ensure predictive accuracy; the output definition, data partition, error measure, and timing metric must remain consistent with the parent physical problem.

Static mechanics and mechanical reduction

The static mechanical study provides a consistent treatment of monolithic and panel-partitioned aligned-orthotropic plates on a regularized unilateral Winkler foundation. The symmetric $C^0$ interior-penalty formulation keeps numerical stabilization inside each panel distinct from the displacement and normal-rotation terms acting across physical seams, while the nonlinear foundation response is resolved by an active-set Newton method. For the monolithic benchmark, the minimum displacement differs from the Navier-series and independent three-dimensional references by approximately 0.05% and 0.15%, respectively. The refinement studies recover the expected polynomial-order hierarchy, and the panel-to-monolithic comparisons show that the finite-penalty discrepancy remains concentrated near the introduced seams and grows in a controlled manner as the number of interfaces increases.

The displacement snapshots exhibit rapid POD spectral decay across the monolithic and panel-partitioned configurations, making POD a common representation for comparing intrusive and non-intrusive reduction within the same mechanical problem. POD–Galerkin improves systematically with basis enrichment, but its online realization still assembles the full-order residual and Jacobian and therefore achieves only modest acceleration. The non-intrusive models replace that online assembly by prediction of the POD coordinates. Across the mechanical benchmarks, POD–GPR provides the strongest overall predictive accuracy, while PODI–RBF gives a favorable accuracy–cost compromise and PODI–Linear and POD–NN provide similarly inexpensive online alternatives with larger errors for the sampled datasets. The predictive surrogates reduce the online cost by approximately three to four orders of magnitude, which is the principal computational benefit for repeated many-query evaluation. Together, these results provide a consistent mechanics-to-ROM workflow in which one parameterized plate model, one displacement output, and one error definition support high-fidelity verification, intrusive reduction, and predictive assessment on freshly generated unseen states.

Coupled thermomechanics and field-wise reduction

The thermomechanical extension connects steady three-dimensional heat transfer to plate bending through displacement-dependent lower-face exchange and the centered thermal driver $\theta$, measured in $\mathrm{K\,m^{-1}}$. A relaxed partitioned iteration alternates the thermal solve, thermal-to-plate reduction, mechanical solve, and displacement update until the coupled state is accepted. The auxiliary analytical comparison, plate-only refinement, and controlled coupled refinement support the thermal, transfer, and mechanical stages. For the reported $(p, p_T, p_\theta) = (2, 1, 2)$ sequences, the isotropic and orthotropic displacement rates are 2.05

and 2.02, and all twelve coupled runs terminate after 14 iterations. The measured cost decomposition further shows that the three-dimensional thermal solve together with thermal-to-plate reduction accounts for 95.6% of the recurrent iteration time on the stated platform, identifying the thermal stage as the principal computational cost.

The reduced-order assessment uses the same accepted coupled states to construct two field-specific workflows: displacement with $N_w = 6$ and thermal driver with $N_\theta = 4$. Each field retains its own output space, metric, basis, physical units, predictor, and error measure, while the sample assignments and full-order provenance remain aligned. The hybrid POD–Galerkin workflow projects the mechanical correction while retaining the thermal solve, thermal-to-plate transfer, and full-order assembly as a coupled projection reference. Among the non-intrusive models, POD–GPR gives the lowest mean thermal-driver error in all three cases and the lowest displacement error in Case 2, whereas direct DL-ROM gives the lowest displacement errors in Cases 3 and 5. These results support a field-wise reduction strategy in which model selection follows the structure of the retained field and its parameter-to-output map.

Transient dynamics and time-conditioned reduction

The transient study extends the plate hierarchy through transverse translational inertia, time-dependent loading, optional algebraic Rayleigh damping, a consistent mass matrix, and average-acceleration Newmark integration. Each trajectory is associated with one prescribed effective foundation regime, and a single active-degree restriction is applied consistently to the mass, stiffness, damping, loading, modal, and reconstructed quantities of panelized systems. Complementary frequency, time-domain, direct-to-modal, energy, harmonic-response, localized-load, and panel-consistency checks support the implementation at the stated discretization. Representative measures include a maximum first-eight-frequency discrepancy of $7.346886 \times 10^{-3}$, direct-to-modal agreement within $8.470401 \times 10^{-7}$, and undamped energy drift of $1.039256 \times 10^{-10}$.

At reduced order, each trajectory contains 41 stored times, complete trajectories are assigned to disjoint training, validation, and test sets before flattening, and only training trajectories enter POD construction, preprocessing, and model fitting. The working dimensions are $N = 7$, 8, and 8 for the three campaigns, with displacement as the reduced output and time as an independent query coordinate. PODI–Linear and POD–GPR reproduce the two load-amplitude campaigns close to the projection benchmark. The varying-stiffness campaign is more demanding because stiffness changes the modal spectrum and accumulated phase; POD–DL-ROM gives the lowest predictive mean, $1.674858 \times 10^{-2}$, while the projection benchmark is $1.615441 \times 10^{-3}$. The comparison therefore separates the quality of the retained spatial representation from the complexity of the parameter–time coordinate map and provides a trajectory-safe family of time-conditioned, non-causal displacement predictors for direct field queries. The reported online timing retains this single-state query interpretation.

Limitations and future outlook

The conclusions arise within the deliberately controlled scope adopted throughout the thesis: small-deflection Kirchhoff–Love kinematics, local foundation and seam idealizations, prescribed parameter domains, and declared output fields. This scope makes it possible to examine the physical models, discretizations, and reduced-order methods through a consistent numerical evidence chain. A natural mechanical extension is to consider moderate and large deformations, geometrically nonlinear plate formulations, and Reissner–Mindlin or higher-order shear-deformation

theories when transverse shear is significant. Material behavior could also be extended beyond linear elasticity to include plasticity, viscoelasticity, damage, and other constitutive nonlinearities, together with more detailed contact, friction, foundation, and panel-connector models.

Several computational developments follow directly from the present results. Hyper-reduction can strengthen the online efficiency of the intrusive mechanical and hybrid thermomechanical models, particularly for nonlinear material, contact, and interface operators. Within the mechanical non-intrusive branch, piecewise-linear POD interpolation is tied to the convex hull of the training parameters, Gaussian-process training becomes increasingly expensive as the database grows, and a single global parameter-to-coordinate map can become less effective as the parameter dimension or response complexity increases. Adaptive sampling, local reduced spaces, and regime-dependent surrogates provide natural extensions, while geometry-aware mappings or local bases are needed when information is to be transferred across substantially different panel layouts. Reduced three-dimensional thermal states and reduced transfer operators can extend the thermomechanical hierarchy beyond the current field-wise outputs. Dedicated output maps can incorporate interface jumps, bending moments, reactions, stresses, and other engineering observables. For transient response, causal structure-preserving reduced dynamics, longer-horizon campaigns, broader loading and damping families, and matched trajectory-level timing can complement the time-conditioned surrogates. Multi-source error and uncertainty indicators can further guide new high-fidelity sampling and distinguish discretization, transfer, compression, regression, and hyper-reduction effects.

Material, component, foundation, connector, and assembled-system experiments would provide the calibration and validation data needed to connect the numerical hierarchy to physical plate systems. Such campaigns should compare model outputs with repeatable measurements across representative support, loading, thermal, and dynamic conditions and should quantify variability that is absent from the present controlled numerical studies.

The verified study-specific ROMs also provide a possible basis for future digital-twin development. Such an extension would require integration with physical sensing, experimentally supported observation and discrepancy models, parameter-identifiability analysis, and validation on instrumented plate systems; these tasks lie beyond the present dissertation.

Future computational studies should preserve versioned environments, meshes, statistical partitions, random seeds, reduced bases, trained models, timing definitions, and checksums so that the physical and computational evidence remains traceable.

## Summary of research contributions

**Static mechanics and mechanical model reduction (Ch. 3–5).**

**Objective:** Develop a high-fidelity and reduced-order framework for repeated analysis of parameterized monolithic and panel-partitioned aligned-orthotropic Kirchhoff–Love plates on a unilateral elastic foundation.

**Method:** Combine the regularized unilateral Winkler law, weak transmission of displacement and normal rotation across physical seams, a symmetric $C^0$ interior-penalty discretization, and an active-set Newton solver; compare intrusive POD–Galerkin with PODI–RBF, PODI–Linear, POD–GPR, and POD–NN for displacement prediction.

**Evidence:** Use Navier-series and independent three-dimensional comparisons, dual-reference refinement, comparisons of panelized and monolithic configurations, and assessment on freshly generated unseen full-order states to support the full-order and reduced-order stages.

**Contribution:** A panel-aware mechanics-to-ROM workflow that preserves the distinction between numerical stabilization and physical seam transmission while carrying one consistent displacement output from verification to predictive many-query assessment.

**Coupled thermomechanics and field-wise model reduction (Ch. 6–8).**

**Objective:** Represent the two-way interaction between steady three-dimensional heat transfer and plate bending and provide efficient forecasts of plate displacement and the thermal driver.

**Method:** Use a relaxed partitioned coupled solve that alternates the three-dimensional thermal solve, thermal-to-plate reduction, mechanical solve, and relaxation; construct separate displacement and thermal-driver one-field ROM workflows with $N_w = 6$ and $N_\theta = 4$ from the same accepted coupled states.

**Evidence:** Use an auxiliary analytical comparison, plate-only and coupled refinement, iteration and cost measurements, and held-out one-field ROM assessments to evaluate the coupled formulation and its field-wise reductions.

**Contribution:** An output-aware reduction design that connects one physically coupled full-order model to separate displacement and thermal-driver ROMs while preserving field-specific spaces, units, metrics, bases, predictors, and error measures.

**Transient dynamics and time-conditioned model reduction (Ch. 9–11).**

**Objective:** Develop high-fidelity simulation and efficient many-query prediction of transient displacement trajectories for parameterized monolithic and panel-partitioned plates.

**Method:** Combine transverse translational inertia, a consistent mass matrix, one active-degree restriction, and average-acceleration Newmark integration; separate complete trajectories before flattening, construct POD and fit the predictors from training trajectories, retain dimensions $N = 7, 8, 8$, and use time as an independent query coordinate.

**Evidence:** Use natural-frequency and analytical time-domain comparisons, direct-to-modal agreement, energy conservation, harmonic and localized-load responses, and panel-consistency checks for the full-order model, together with held-out complete-trajectory ROM assessment.

**Contribution:** A trajectory-safe, time-conditioned displacement-ROM workflow that links verified full-order trajectories to held-out field prediction through consistent data partitioning, training-based POD construction, and cross-campaign assessment.

**Thesis-level integration of the three model–FOM–ROM sequences.** Taken together, the three physical studies form one cumulative dissertation architecture while retaining their distinct states, assumptions, outputs, and evidence. They reuse a common vocabulary for the plate, compliant support, panel partition, and $C^0$ interior-penalty approximation; preserve the distinction between numerical facet stabilization and physical seam transmission; and employ study-specific output maps, data partitions, POD spaces, predictors, error measures, and timing protocols. Across the mechanical, thermomechanical, and transient studies, complementary full-order verification or solver-assessment programs are combined with projection diagnostics and held-out predictive tests. This integration provides a coherent and traceable model–FOM–ROM hierarchy for parameterized orthotropic plate systems, supporting output-specific many-query prediction without conflating distinct physical regimes or assessment claims.

Final perspective

The dissertation shows how parameterized plate models can be carried through a disciplined computational chain: explicit physical assumptions, panel-aware high-fidelity discretization, numerical verification, declared output maps, output-specific reduction, and held-out many-query assessment. The value of the hierarchy lies in preserving traceability across these stages while allowing each physical study to retain its own state, data, and evidence. The resulting models and workflows provide an efficient basis for repeated structural prediction within their stated domains and a coherent foundation for future experimental calibration and instrumented extensions.

# Computational Implementation and Reproducibility

# A


A.1 Purpose and scope . . . . 146
A.2 Software, artifacts, and shared finite element architecture . . . . . 146
A.3 Study-specific full-order workflows . . . . . 148
A.4 Snapshot, output-transfer, and reduced-order pipelines . . . . . 150
A.5 Reduced-order model architectures and hyperparameters . . . . . 152
A.6 Code and computational reproducibility . . . . . 156


## A.1 Purpose and scope

This appendix documents the computational implementation supporting the numerical studies in the preceding chapters. It records the software stack, finite element workflows, output transfers, data partitions, reduced-model artifacts, post-processing conventions, and timing definitions needed to trace a reported result to its computational origin [108, 109].

The implementation hierarchy is

$$\begin{gathered}\text{parameter record, geometry, mesh, and solver configuration}\\ \downarrow\\ \text{full-order solve, convergence checks, and accepted state}\\ \downarrow\\ \text{declared output transfer, snapshot database, and data partition}\\ \downarrow\\ \text{basis or latent representation, predictive model, and assessed output.}\end{gathered} \tag{A.1}$$

Each stage records the information needed to reproduce its output. A field sample includes the parameter vector, units, discretization, degree-of-freedom ordering, solver status, output-space definition, and statistical-subset membership. A serialized ROM additionally stores its basis, scalers, split indices, hyperparameters, software version, and trained estimator. Numerical verification records the reference construction, error norm, and comparison space. Section A.6 identifies the project repository containing the executable workflows, configurations, and supporting artifacts.

## A.2 Software, artifacts, and shared finite element architecture

Python orchestrates the workflows [110]; UFL forms are assembled with FEniCS/DOLFIN [37, 111], and PETSc is accessed through DOLFIN for sparse linear algebra. NumPy, SciPy, scikit-learn, PyTorch, pandas, and Matplotlib support sampling, array processing, interpolation, regression, neural training, data handling, and visualization [112–117]. POD and reduced-basis utilities also use RBNiCS-related routines [43, 118], while notebooks and reusable Python modules act as controlled experiment drivers [119].

The software versions and hardware configuration used for a reported computation form part of its provenance. For example, the coupled refinement profile in Chapter 7 used Python 3.12.13, DOLFIN 2019.1.0, UFL 2019.1.0, PETSc/petsc4py 3.21.2, and mpi4py 3.1.6 on macOS 26.5.2, with an Apple M5, 16 GB of memory, an x86_64 FEniCS runtime through Rosetta 2, one MPI rank, and one OpenMP/BLAS thread. This record is specific to that profile; the repository associates other reported results

**Table A.1:** Study-specific computational states, active reduced outputs, and principal online operations.

| Workflow | Full-order state | Active reduced output | Principal online operation |
|---|---|---|---|
| Static mechanical | Panel-wise or monolithic displacement and the active compression–uplift partition | Native displacement for intrusive projection or transferred global continuous displacement for non-intrusive prediction | Active-set Newton solve, non-hyper-reduced POD–Galerkin equilibrium, or POD-coordinate prediction |
| Thermomechanical | Three-dimensional temperature, plate thermal driver, displacement, and partitioned-coupling status | Two independent one-field outputs: displacement with $N_w = 6$ and thermal driver with $N_\theta = 4$, each with its own field space, basis, and model | Coupled FOM iteration, hybrid mechanical projection, or one field-specific non-intrusive predictor |
| Transient dynamics | Displacement, velocity, and acceleration histories on the active dynamic degrees of freedom | Global displacement at 41 stored times; final POD dimensions 7, 8, and 8 for the three campaigns | Complete Newmark trajectory or non-causal parameter–time prediction of one displacement state |

with their corresponding environment, hardware, and thread metadata.

Table A.1 separates the complete full-order state from the quantity actually reduced or observed. This is essential because the thermomechanical FOM is physically coupled but its ROMs are field-wise, whereas the transient integrator advances displacement, velocity, and acceleration but reduces displacement only.

The shared plate code starts from a parameter record $\mu$, a rectangular midsurface, and either a monolithic domain or a non-overlapping panel partition. Geometry construction creates cell, exterior-boundary, numerical interior-facet, and physical-interface markers. Two distinct internal-boundary treatments are retained:

- facets strictly inside one physical panel carry the consistency, symmetry, and normal-slope penalty terms of the symmetric $C^0$ interior-penalty discretization;
- physical seams carry separate displacement and normal-rotation transmission penalties.

The first family is numerical stabilization; the second represents the adopted panel-coupling model and is not interpreted as an experimentally calibrated connector law.

A monolithic plate uses one scalar Lagrange space, whereas a panelized plate retains one component space per panel and couples independent traces at seams. The resulting mixed representation contains component degrees of freedom outside their physical panels. They can remain in static residual assembly but have zero mass in the transient problem. Chapter 10 therefore constructs one active index set from the mass diagonal and applies it consistently to mass, stiffness, damping, loads, initial data, eigenvectors, direct and modal histories, lifted states, and snapshots; any inconsistent restriction changes the semidiscrete system or leaves a singular mass matrix. Each stored state is accompanied by parameter-dependent rigidities, support and interface coefficients, loads, thermal data, damping, time grids, mesh identifiers, polynomial degrees, quadrature settings, and solver configuration. Mesh identifiers resolve to immutable generator inputs and checksums rather than machine-specific absolute paths.

## A.3 Study-specific full-order workflows

**Static mechanical workflow.** The mechanical solver implements Chapters 3 and 4. For each parameter instance it assembles orthotropic bending, panel-interior $C^0$-IPG terms, physical-seam penalties, the regularized unilateral Winkler term, and the signed transverse load. The foundation is the only nonlinear contribution. Compression and uplift regions are inferred from the current iterate, frozen for one branchwise tangent solve, and updated after the Newton correction.

**Mechanical full-order solve**

For one mechanical parameter instance:

1. build the mesh and the cell, exterior-boundary, numerical-facet, and physical-seam markers;
2. assign the orthotropic rigidities, foundation data, interface weights, load data, polynomial degree, and solver tolerances;
3. assemble the nonlinear residual and its active-branch Jacobian;
4. solve the Newton correction, update the displacement, and recompute the compression–uplift partition;
5. accept the state only when the declared algebraic residual criterion is satisfied;
6. store the displacement, active-set status, parameters, discretization, solver metadata, and any derived moment fields.

For algebraic displacement $\mathbf{w}_h^{(k)}$ and residual $\mathbf{F}_h(\mathbf{w}_h^{(k)})$, the stopping test is

$$\left\|\mathbf{F}_h(\mathbf{w}_h^{(k)})\right\|_{\ell^2} \le \max\left\{\varepsilon_{\mathrm{abs}}, \varepsilon_{\mathrm{rel}}\left\|\mathbf{F}_h(\mathbf{w}_h^{(0)})\right\|_{\ell^2}\right\}. \tag{A.2}$$

The active set accepted under Equation (A.2) is induced by the current iterate. Numerical tolerance values belong to the campaign configuration and are not inferred from the symbolic criterion. Moment resultants are post-processed from the accepted Kirchhoff–Love displacement and are not additional unknowns.

Reproducing the dual-reference and refinement evidence also requires the Navier series settings, independent three-dimensional solid configuration, fine-mesh reference identifiers, transfer rules, and exact $L^2(\Omega)$ and $H^1(\Omega)$ error definitions. Panel studies additionally require the partition, seam markers, degree-dependent penalties, and monolithic comparison state.

**Thermomechanical workflow.** The implementation follows Chapters 6 and 7. A tetrahedral mesh resolves the three-dimensional thermal body and an independent panel-aligned triangular mesh resolves the plate. For prescribed displacement, the thermal solve includes conduction, top-face convection, optional long-wave radiation, absorbed solar flux, and lower-face substrate exchange; the exchange coefficient is evaluated by mapping each thermal quadrature point to the midsurface and sampling the current displacement. Gauss–Legendre thickness quadrature then forms the centered first moment and interpolates it to the plate thermal-driver space. Thus $\theta_h$, measured in $\mathrm{K\,m^{-1}}$, is not a temperature. The mechanical solve uses $\theta_h$ through the thermal bending functional while retaining the unilateral foundation and seam terms. The primary FOM is a relaxed partitioned fixed point, with relaxation applied only to displacement.

**Partitioned thermomechanical solve**

Given $w_h^{(0)}$, $0 < \omega \le 1$, tolerances $\mathrm{tol}_w$ and $\mathrm{tol}_\theta$, a safeguard $\varepsilon_0 > 0$, and $K_{\max}$:

1. solve the three-dimensional thermal problem using $w_h^{(k)}$ in the lower-face exchange law;
2. reconstruct $\theta_h^{(k+1)}$ from the temperature increment;
3. solve the nonlinear plate problem with $\theta_h^{(k+1)}$ fixed, producing $\widetilde{w}_h^{(k+1)}$;
4. apply
$$w_h^{(k+1)} = (1-\omega) w_h^{(k)} + \omega \widetilde{w}_h^{(k+1)};$$
5. after relaxation, evaluate the relative finite element $L^2(\Omega)$ updates of $w_h$ and $\theta_h$;
6. accept the coupled state only when both declared tolerances are satisfied; otherwise continue to $K_{\max}$ and mark the run nonconverged.

At termination, the returned thermal field and driver were generated from $w_h^{(k)}$, while the accepted displacement is the relaxed $w_h^{(k+1)}$. The residual mismatch is controlled by the update tolerances. A minimum coupled-state record is

$$\left(\boldsymbol{\mu}, T_h, \theta_h, w_h, \omega, \mathrm{tol}_w, \mathrm{tol}_\theta, K_{\max}, K_{\mathrm{used}}, e_w, e_\theta, \mathrm{status}\right), \tag{A.3}$$

supplemented by thermal and mechanical Newton tolerances, linear solvers, preconditioners, mesh and quadrature identifiers, and transfer metadata. A one-way evaluation is the explicit specialization $K_{\max} = 1$ and $\omega = 1$; its displacement is not returned to a second thermal solve.

The final one-field ROM databases use

$$N_\Omega = 64, \qquad p = 2, \qquad (N_x^T, N_y^T, N_z^T) = (64, 32, 16), \qquad p_T = p_\theta = 1,$$

with 20 through-thickness quadrature points, $\omega = 0.7$, $\mathrm{tol}_w = \mathrm{tol}_\theta = 10^{-5}$, and $K_{\max} = 25$. These data-generation settings differ from the tighter, higher-order verification configurations of Chapter 7.

**Transient dynamics workflow.** The transient solver implements Chapters 9 and 10 under the assumptions stated in Section 9.1. One prescribed effective Winkler coefficient is fixed for each trajectory, and only transverse translational inertia enters the governing dynamic model. Rayleigh damping is available at the algebraic level, but the ROM campaigns use zero damping.

After essential-boundary elimination and active-degree restriction,

$$\mathbf{M}_a \ddot{\mathbf{w}}_a(\tau) + \mathbf{C}_a \dot{\mathbf{w}}_a(\tau) + \mathbf{K}_a \mathbf{w}_a(\tau) = \mathbf{f}_a(\tau). \tag{A.4}$$

Average-acceleration Newmark uses $\beta = 1/4$ and $\gamma = 1/2$. The compatible initial acceleration follows from the restricted equilibrium, and a constant effective operator is prepared once and reused when the operator and time step are fixed.

**Transient trajectory generation**

For one physical trajectory parameter:

1. assemble the mixed-space mass, stiffness, damping, and load objects;
2. construct the active indices from the mass diagonal and restrict all matrices, vectors, initial data, and output maps consistently;
3. compute the compatible initial acceleration and form the constant Newmark effective operator;
4. advance the complete time grid and retain any verification histories requested by the campaign;
5. lift each stored active displacement to the mixed representation and, for panelized cases, transfer the physical panel field to the declared global continuous output space;
6. store the physical parameters, time grid, support regime, damping, active-index map, solver status, and temporal-storage indices.

The final ROM campaigns use $N_\Omega = 64$, $p = 2$, $T = 0.20\,\mathrm{s}$, $\Delta\tau = 10^{-3}\,\mathrm{s}$, 200 integration steps, storage every five steps, free exterior edges, the prescribed compression coefficient, zero initial displacement and velocity, zero Rayleigh damping, and a $0.351\,\mathrm{m} \times 0.351\,\mathrm{m}$ step-loaded patch. Each trajectory therefore contains 41 stored displacement fields. Velocity, acceleration, stress resultants, reactions, energy, and seam jumps may be retained for FOM verification but are not reduced outputs in Chapter 11.

## A.4 Snapshot, output-transfer, and reduced-order pipelines

An offline campaign evaluates a declared parameter design and stores accepted FOM outputs in deterministic coefficient spaces. For a generic one-field output $\widehat{\chi}_h$,

$$\mathbf{S}_\chi = \left[\widehat{\boldsymbol{\chi}}_h^{(1)} \mid \cdots \mid \widehat{\boldsymbol{\chi}}_h^{(N_s)}\right]. \tag{A.5}$$

The hat permits a transferred rather than native FOM vector. Panel-wise mechanical and transient states are mapped by discrete $L^2(\Omega)$ projection to a fixed continuous space with common coefficient ordering. This does not alter the parent FOM, but it does omit seam jumps as independent ROM outputs; an interface-aware surrogate would require seam observables or the native broken state.

**Mechanical data and reduction.** The six-parameter monolithic campaign contains 1000 Latin-hypercube snapshots; the two-panel and each reported multi-panel campaign contain 300. Final method comparisons use 100 test parameters. Intrusive POD uses the native FOM space, whereas non-intrusive models use the transferred global displacement and the displacement-plus-gradient Gram matrix. Campaign-specific basis and regression indices are preserved with each artifact.

The predictive inventory is PODI–RBF, PODI–Linear, POD–GPR, and POD–NN; POD–Proj is diagnostic. The intrusive POD–Galerkin model evaluates the active set on the lifted reduced state but assembles the

full-order residual and Jacobian before projection, so it is not hyper-reduced. Exact surrogate and network configurations are summarized in Section A.5.

**Thermomechanical field-wise databases and POD.** Each accepted coupled state supplies two distinct outputs,

$$\widehat{w}_h(\boldsymbol{\mu}) \in X_h^w, \qquad \widehat{\theta}_h(\boldsymbol{\mu}) \in X_h^\theta.$$

They share case-level sample assignments but have separate finite element spaces, Gram matrices, POD metrics, bases, scalers, models, errors, and units. Displacement uses the displacement-plus-gradient metric; the thermal driver uses the finite element $L^2(\Omega)$ metric. The three-dimensional temperature remains part of the coupled FOM, not an active ROM output.

The accepted records and source-global partitions are

| case | accepted rows | train/val./test | dim $X_h^w$ | dim $X_h^\theta$ | $N_w$ | $N_\theta$ |
|---|---|---|---|---|---|---|
| 2 | 450 | 360/45/45 | 20,503 | 5,194 | 6 | 4 |
| 3 | 378 | 302/38/38 | 20,375 | 5,160 | 6 | 4 |
| 5 | 500 | 400/50/50 | 20,503 | 5,194 | 6 | 4 |

Seed 100 selects training rows and seed 101 splits the remainder equally into validation and test; both fields use identical source-global indices. Only training rows enter POD and fitting. Validation supports dimension studies, neural early stopping, and checkpoint selection; test rows support final accuracy, timing, and spatial diagnostics.

The field-wise non-intrusive methods are PODI–RBF, PODI–Linear, POD–GPR, POD–NN, POD–DL-ROM, and direct DL-ROM. Direct DL-ROM uses a separate full-state autoencoder/decoder for each field. Hybrid POD–Galerkin reduces only the mechanical correction and retains the three-dimensional heat solve, thermal-to-plate reduction, and full-order mechanical residual and Jacobian.

**Transient complete-trajectory protocol.** The databases preserve trajectory and stored-time axes before flattening. They contain 200 Case 2 monolithic, 250 two-panel Case 2, and 250 Case 3 monolithic trajectories, each with 41 stored times and continuous-output dimensions 20,503, 20,375, and 20,503. The two-panel realization has 40,750 mixed-space degrees of freedom before filtering and 20,488 active dynamic unknowns. Seed 100 assigns complete trajectories to the 80% training subset; seed 101 splits the remainder equally between validation and test:

| campaign | trajectory split | row split |
|---|---|---|
| Case 2 monolithic | 160/20/20 | 6560/820/820 |
| two-panel Case 2 | 200/25/25 | 8200/1025/1025 |
| Case 3 monolithic | 200/25/25 | 8200/1025/1025. |

All 41 states of one trajectory inherit the same assignment; validation and test trajectories enter neither compression nor fitting.

Training-only uncentered snapshot rows first undergo a Euclidean SVD. The spatial vectors are then converted to finite element functions and reorthonormalized in the declared displacement $H^1$-type product before coordinates and projection diagnostics are computed. This is Euclidean

spectral extraction followed by finite element reorthonormalization, not a purely $H^1$-optimal POD. The retained dimensions are $N = 7$ for Case 2 monolithic and $N = 8$ for the two-panel Case 2 and Case 3 monolithic campaigns.

The predictors are PODI–RBF, PODI–Linear, POD–GPR, POD–NN, and a non-causal, time-conditioned POD–DL-ROM. Its coefficient autoencoder is auxiliary, the unequal executed hidden widths instantiate no residual skip, and the known zero initial displacement is imposed at prediction time. The complete model definitions appear in Section A.5.

For every POD-based method, let $\mathbf{a}_N$ denote the local reduced coordinate vector. Reconstruction is

$$\widehat{\boldsymbol{\chi}}_{h,N}(\boldsymbol{\xi}) = \overline{\boldsymbol{\chi}} + \mathbf{V}_N \mathbf{a}_N(\boldsymbol{\xi}), \tag{A.6}$$

where the mean is retained only for a centered source artifact, $\boldsymbol{\xi} = \boldsymbol{\mu}$ for static or steady outputs, and $\boldsymbol{\xi} = (\boldsymbol{\mu}, \tau)$ for transient displacement. Time is an evolution variable and query coordinate, not a physical parameter; the transient surrogate predicts the requested field directly and does not advance a reduced Newmark state.

Each serialized predictor binds its basis version, field coordinates, input and output transforms, hyperparameters, seeds, data indices, estimator parameters or weights, and software environment. Direct DL-ROM additionally requires its encoder/decoder and output scaling. POD–Proj remains a compression diagnostic because it starts from an available FOM state. The corresponding basis, preprocessing, estimator, and evaluation files are organized in the repository described in Section A.6.

## A.5 Reduced-order model architectures and hyperparameters

This section complements the data-lineage and POD definitions in Section A.4 and records the architectures used in Chapters 5, 8, and 11. For one output $\chi \in \{w, \theta\}$, the current uncentered basis $\mathbf{V}_N^{\chi}$ and metric $\mathbf{G}_{\chi}$ define training coordinates $\mathbf{a}_N^{(i)} = (\mathbf{V}_N^{\chi})^{\top} \mathbf{G}_{\chi} \widehat{\boldsymbol{\chi}}_h^{(i)}$. Inputs and coordinates are standardized from training data only. Static and steady models query $\boldsymbol{\mu}$; transient models query $(\boldsymbol{\mu}, \tau)$. Thermomechanical displacement and thermal-driver models remain independent with $N_w = 6$ and $N_\theta = 4$, while the three transient campaigns use $N = 7, 8, 8$.

**Common projection and surrogate architecture.** Figure A.1 shows the shared POD path. POD–Galerkin obtains coordinates from a projected equilibrium or correction solve; POD–Proj projects an available FOM field and is diagnostic; PODI–Linear, PODI–RBF, and POD–GPR learn the query-to-coordinate map and share the same inverse scaling and POD lifting.

**POD–NN architecture.** The shared POD–NN is a fully connected coordinate regressor with three ELU hidden layers of widths (128, 96, 64), followed by a linear $N_\chi$-dimensional coordinate output. All three studies use coefficient-only training on precomputed standardized POD projection coordinates. Figure A.2 separates the training targets from the

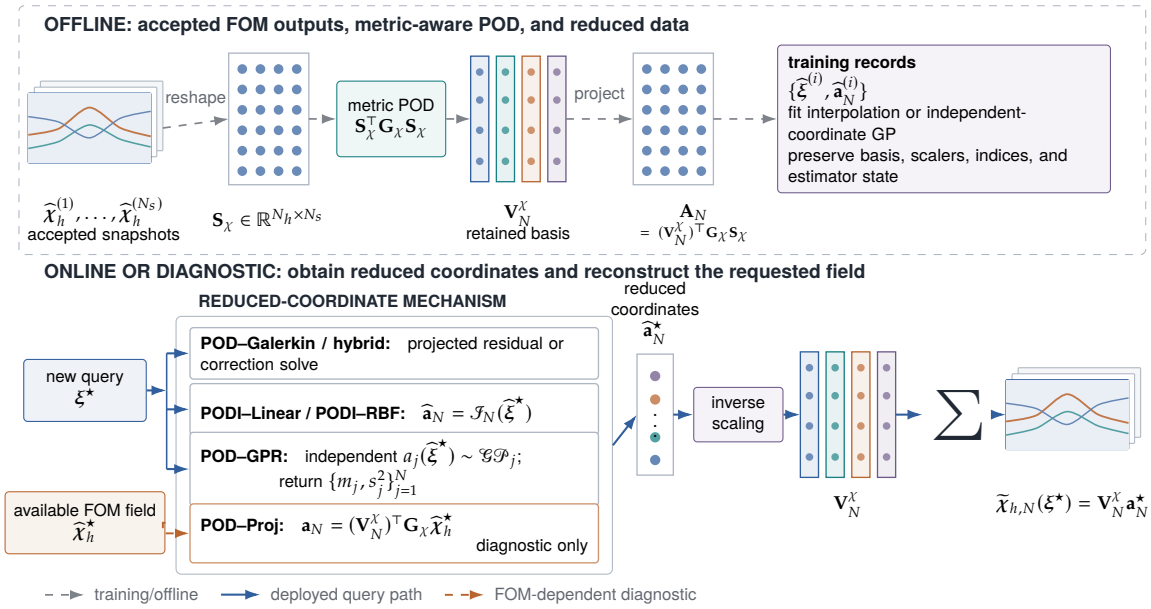



**Figure A.1:** Common POD projection, interpolation, and regression architecture used in the thesis. The upper band constructs the field-specific basis and reduced database. The lower band distinguishes projected residual or correction solves, deterministic POD interpolation, independent-coordinate POD–GPR, and the non-predictive POD–Proj diagnostic. All predictive paths terminate in the same inverse scaling and POD reconstruction.

deployed query path, in which inverse coordinate scaling and POD lifting follow coordinate prediction.

**POD–DL-ROM and direct DL-ROM architectures.** The steady POD–DL-ROM acts in coefficient space: its encoder is $N_\chi \to 128 \to 96 \to 64 \to n_z$, its mirrored decoder is $n_z \to 64 \to 96 \to 128 \to N_\chi$, and its parameter map is $P \to 128 \to 96 \to 64 \to n_z$, with $n_z = N_\chi$. Direct DL-ROM uses field-specific full-state autoencoders: displacement has widths $(2048, 1024, 512, 256)$ with $n_z^w = 6$, while the thermal driver has $(1024, 512, 256, 128)$ with $n_z^\theta = 4$. Both use ELU hidden activations with linear latent and output layers. The transient POD–DL-ROM uses enriched parameter–time features and $(192, 160, 128)$ SiLU layers with layer normalization and dropout 0.02; its feature-to-coordinate map is the online predictor, while the coefficient autoencoder supports training diagnostics. Figure A.3 records the three executed paths.

**Principal implementation settings.** Table A.2 summarizes the six predictive ROM families by their executed architecture or algorithm and the common implementation used across the corresponding studies. Detailed training schedules and case-specific assessment settings are reported in the relevant chapters.

**ROM post-processing and timing.** Post-processing produces displacement and thermal-driver fields, moment and absolute-error maps, probe and line histories, refinement curves, POD spectra, and coordinate diagnostics. A reproducible figure requires the underlying arrays, sample or trajectory identifier, ordering rule, units, color limits, and generation script; a PDF or PNG preserves appearance but not numerical provenance.

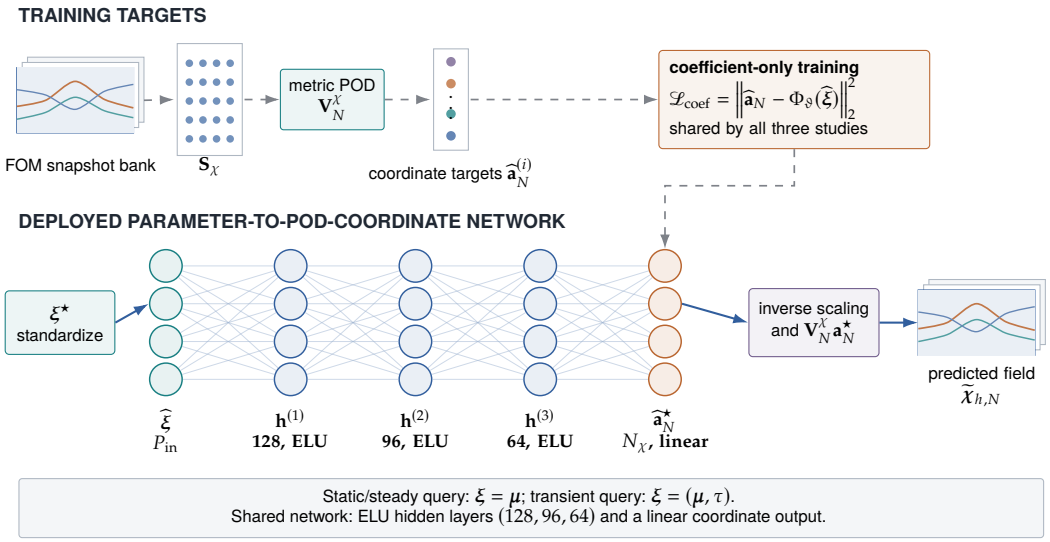



**Figure A.2:** Implemented coefficient-only POD–NN architecture. Training fields are compressed by the field-specific POD basis to form standardized coordinate targets. For the mechanical, thermomechanical, and transient studies, the network minimizes only the POD-coordinate loss. The deployed network maps the standardized physical query through the exact hidden widths (128, 96, 64), with ELU activations and a linear coordinate output; inverse scaling and POD lifting then recover the displacement or thermal-driver field. Dashed connectors denote training-only target construction and coefficient-loss evaluation.

**Table A.2:** Principal architecture or algorithm and common implementation of the six predictive ROM families.

| Method | Executed architecture or algorithm | Common implementation |
|---|---|---|
| PODI–RBF | Vector-valued thin-plate-spline interpolation of standardized POD coordinates, with $\varphi(\rho) = \rho^2 \log \rho$ and affine augmentation $\mathbf{b}_0 + \mathbf{B}\widehat{\boldsymbol{\xi}}$. | Standardized inputs and coordinates, standard moment constraints, and POD reconstruction after coordinate prediction. Smoothing and neighborhood settings follow the study-specific configuration. |
| PODI–Linear | Delaunay tessellation in standardized input space with barycentric interpolation on the containing simplex. | Interpolation is restricted to the training convex hull. Exterior or non-finite queries use a deterministic nearest-training-neighbor fallback. |
| POD–GPR | One independent scalar Gaussian process per standardized retained POD coordinate; posterior means are POD-lifted and marginal coordinate variances are retained. | Study-specific covariance kernel, optional white-noise term, separate numerical nugget, multiple optimizer starts, and training-only fitting. |
| POD–NN | $P_{\text{in}} \to 128 \to 96 \to 64 \to N_\chi$; three ELU hidden layers and a linear coordinate output. | Coefficient-only MSE on standardized POD projection coordinates, with AdamW, validation-based scheduling and early stopping, best-state restoration, inverse scaling, and POD lifting. |
| POD–DL-ROM | Coefficient autoencoder $N_\chi \to 128 \to 96 \to 64 \to n_z \to 64 \to 96 \to 128 \to N_\chi$, with parameter map $P \to 128 \to 96 \to 64 \to n_z$. The time-conditioned regressor uses $P_{\text{feat}} \to 192 \to 160 \to 128 \to N$. | Latent or coordinate learning followed by decoder-based field reconstruction. The time-conditioned variant uses enriched inputs, SiLU activations, layer normalization, and dropout 0.02. |
| DL-ROM | Direct full-field autoencoders: $N_h^w \to 2048 \to 1024 \to 512 \to 256 \to 6$ and $N_h^\theta \to 1024 \to 512 \to 256 \to 128 \to 4$, with parameter map $P \to 128 \to 96 \to 64 \to n_z$. | Direct full-field latent reconstruction with ELU hidden activations, linear latent and output layers, and field-specific scalers and checkpoints. |

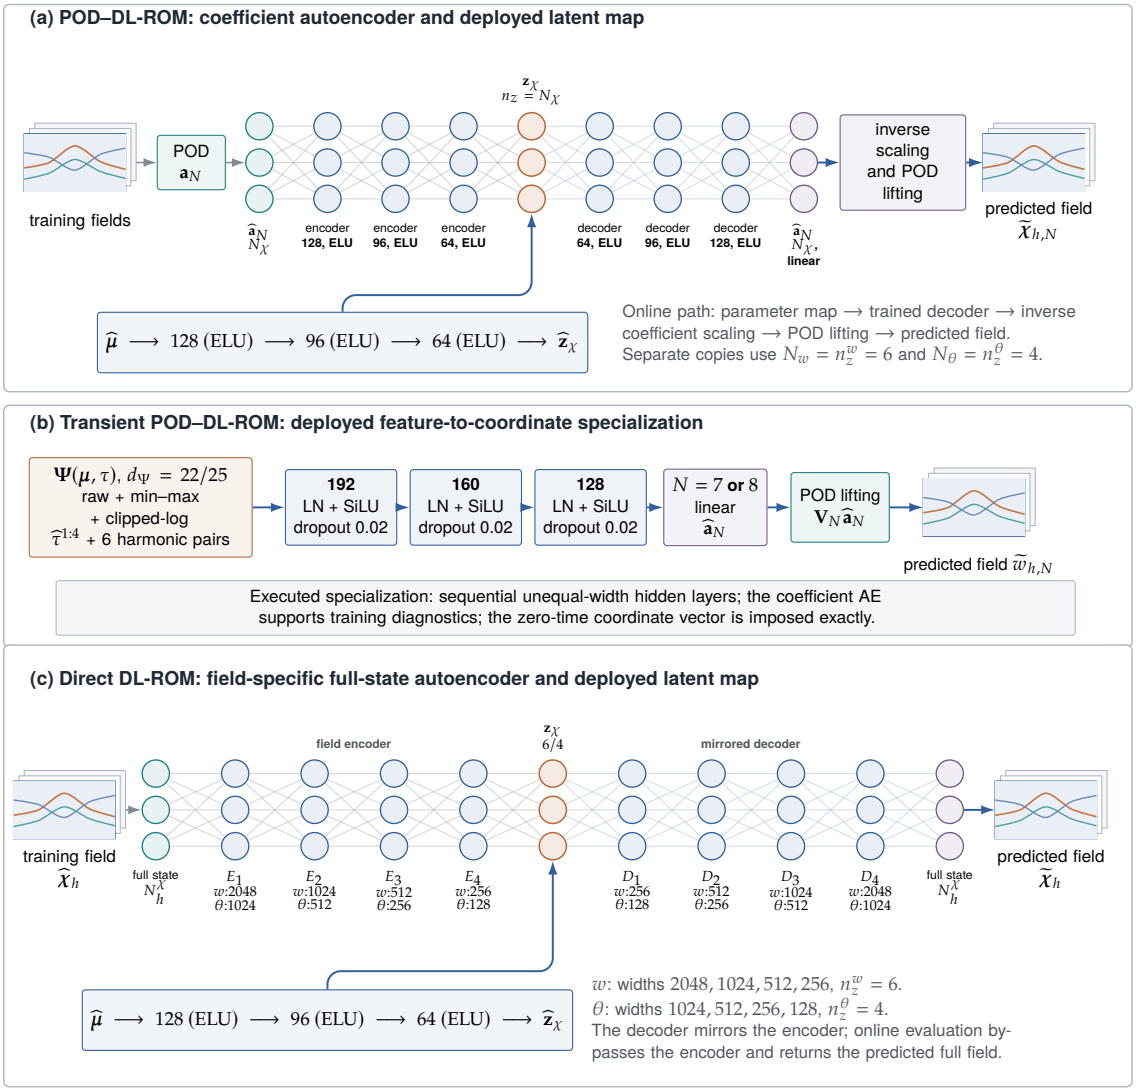



**Figure A.3:** Implemented deep-learning ROM architectures. Panel (a) shows the coefficient-space POD–DL-ROM autoencoder and the deployed parameter-to-latent-to-decoder path, including inverse scaling, POD lifting, and the predicted field. Panel (b) records the distinct time-conditioned transient POD–DL-ROM, whose enriched feature network maps directly to displacement POD coordinates and then to the predicted displacement field. Panel (c) shows the independent displacement and thermal-driver direct DL-ROMs, including their full-state decoder outputs; paired encoder/decoder widths refer to $w/\theta$, respectively. Gray dashed connectors are training-only, whereas blue connectors belong to deployed online paths.

For a ROM output $\chi \in \{w, \theta\}$ in a declared finite element space,

$$e_{L^2,\chi} = \frac{\|\chi_{\mathrm{FOM}} - \chi_{\mathrm{ROM}}\|_{L^2(\Omega)}}{\|\chi_{\mathrm{FOM}}\|_{L^2(\Omega)}}. \tag{A.7}$$

Both fields are transferred to the same output space before Equation (A.7) is evaluated. Thermomechanical displacement and thermal-driver errors use their own mass matrices. Transient errors use the continuous-displacement mass matrix and exclude the exact zero-time state from relative summaries because its denominator vanishes.

Timing records identify the measured operation and its inclusion scope. For the refined coupled-FOM profile in Chapter 7, all 12 retained runs terminate after 14 partitioned iterations. Median recurrent time is split among the three-dimensional thermal PDE (85.4%), thermal-to-plate reduction and transfer (10.2%), and the mechanical PDE (4.4%); the first two sum to 95.6%. Complete profiled wall times are 3199.68 s and

4978.89 s for the isotropic and orthotropic refined cases on the recorded platform. These values describe the chapter-specific setup, iteration, and finalization operations on that platform.

The ROM timing conventions are:

- mechanical and thermomechanical predictive-ROM ratios compare matched single-query FOM and reconstructed one-field ROM evaluations under the chapter-specific online definitions; offline generation, basis construction, training, model loading, and plotting are excluded;
- POD–Proj timing reports the projection of an already available FOM state and is interpreted as a compression-cost diagnostic;
- the transient ratio compares one complete FOM trajectory with the prediction and reconstruction of one ROM state and is reported explicitly as a trajectory-to-single-state cost ratio.

## A.6 Code and computational reproducibility

The project repository for the computational implementation and supporting research artifacts developed for this dissertation is

`parametric-plate-rom`.

The repository serves as the executable companion to Chapters 3–11 and to this appendix. It contains the static mechanical, thermomechanical, and transient full-order solvers and verification cases; mesh-generation and field-transfer utilities; snapshot and complete-trajectory data pipelines; POD basis construction; intrusive, hybrid, non-intrusive, and deep-learning ROM training and assessment; and the scripts used for post-processing and figure generation. The corresponding configuration files, software-environment specifications, and execution instructions are provided with the relevant workflows.

## Use of artificial intelligence tools

Generative artificial intelligence tools were used as assistive aids in debugging and refactoring code and in improving language and LaTeX typesetting. Their outputs were treated as suggestions and were critically reviewed, tested, corrected, and integrated by the author prior to use. The mathematical models, derivations, scientific content, computational methods, parameter choices, numerical experiments, software implementations, verification studies, interpretation of results, and conclusions remain the responsibility of the author. No numerical data or simulation results were generated by an artificial intelligence system; all reported computational evidence was produced through executable workflows run and verified by the author.

# Bibliography


[1] Eugene Boadi-Danquah et al. "Lightweight modular steel floor system for rapidly constructible and reconfigurable buildings." In: *International Journal of Computational Methods and Experimental Measurements* 5.4 (2017), pp. 562–573. DOI: 10.2495/CMEM-V5-N4-562-573 (cit. on p. 2).

[2] Dina Kuttah. "A review on temporary roads and portable matting systems improving traffic resilience and disaster mobility." In: *Journal of Disaster Science and Management* 2 (2026), p. 5. DOI: 10.1007/s44367-026-00030-w (cit. on p. 2).

[3] Gianluigi Rozza, Dinh Bao Phuong Huynh, and Anthony T. Patera. "Reduced Basis Approximation and A Posteriori Error Estimation for Affinely Parametrized Elliptic Coercive Partial Differential Equations: Application to Transport and Continuum Mechanics." In: *Archives of Computational Methods in Engineering* 15.3 (2008), pp. 229–275. DOI: 10.1007/s11831-008-9019-9 (cit. on p. 2).

[4] Alfio Quarteroni, Gianluigi Rozza, and Andrea Manzoni. "Certified Reduced Basis Approximation for Parametrized Partial Differential Equations and Applications." In: *Journal of Mathematics in Industry* 1.1, 3 (2011). DOI: 10.1186/2190-5983-1-3 (cit. on p. 2).

[5] EPS. *Drivable Event Flooring: For Construction Sites and Events*. 2026. URL: https://eps.net/en/products/drivable-event-flooring/ (visited on 05/16/2026) (cit. on p. 2).

[6] EPS. *Pedestrian Event Flooring*. 2026. URL: https://eps.net/en/products/pedestrian-event-flooring/ (visited on 05/16/2026) (cit. on p. 2).

[7] EPS. *Temporary Event Flooring*. 2026. URL: https://eps.net/en/industries/festivals/event-flooring/temporary-event-flooring/ (visited on 05/16/2026) (cit. on p. 2).

[8] EPS. *Arena Panels*. 2026. URL: https://eps.net/en/products/drivable-event-flooring/arena-panels/#&gid=1&pid=3 (visited on 05/09/2026) (cit. on pp. 2, 3).

[9] EPS. *Armordeck 3*. 2026. URL: https://eps.net/en/products/drivable-event-flooring/armordeck-3/#&gid=1&pid=2 (visited on 05/09/2026) (cit. on p. 3).

[10] EPS. *Arena Light Panels*. 2026. URL: https://eps.net/en/products/drivable-event-flooring/arena-light-panels/#&gid=1&pid=1 (visited on 05/09/2026) (cit. on p. 3).

[11] S. Timoshenko and J.N. Goodier. *Theory of Elasticity*. 2nd ed. McGraw-Hill Book Company, 1951 (cit. on pp. 2, 35).

[12] S. Timoshenko and S. Woinowsky-Krieger. *Theory of Plates and Shells*. 2nd ed. New York: McGraw-Hill, 1959 (cit. on pp. 2, 10, 11, 14, 21, 34, 104).

[13] J.N. Reddy. *Theory and Analysis of Elastic Plates and Shells*. 2nd ed. CRC Press, 2007 (cit. on pp. 2, 10–12, 14, 21, 34, 35, 66, 79, 104, 115, 116).

[14] E. Ventsel and T. Krauthammer. *Thin Plates and Shells: Theory, Analysis, and Applications*. New York: Marcel Dekker, 2001 (cit. on pp. 2, 10–12, 14, 21, 104).

[15] A. D. Kerr. "Elastic and Viscoelastic Foundation Models." In: *Journal of Applied Mechanics* 31.3 (1964), pp. 491–498. DOI: 10.1115/1.3629667 (cit. on pp. 3, 12, 105).

[16] Y. Weitsman. "On Foundations That React in Compression Only." In: *Journal of Applied Mechanics* 37.4 (1970), pp. 1019–1030. DOI: 10.1115/1.3408653 (cit. on pp. 3, 12).

[17] Michael F. Modest and Sandip Mazumder. *Radiative Heat Transfer*. 4th ed. Academic Press, 2021 (cit. on pp. 4, 15, 64).

[18] John A. Duffie and William A. Beckman. *Solar Engineering of Thermal Processes*. 4th ed. Hoboken, NJ: John Wiley & Sons, 2013 (cit. on pp. 4, 15, 64).

[19] Chakravarti V. Madhusudana. "Gap Conductance at the Interface." In: *Thermal Contact Conductance*. 2nd ed. Mechanical Engineering Series. Cham: Springer, 2014, pp. 55–77. DOI: 10.1007/978-3-319-01276-6_4 (cit. on pp. 4, 15, 65).

[20] M. G. Cooper, B. B. Mikic, and M. M. Yovanovich. "Thermal Contact Conductance." In: *International Journal of Heat and Mass Transfer* 12.3 (1969), pp. 279–300. DOI: 10.1016/0017-9310(69)90011-8 (cit. on pp. 4, 15, 65).

[21] Carlos A. Felippa, K. C. Park, and Charbel Farhat. "Partitioned Analysis of Coupled Mechanical Systems." In: *Computer Methods in Applied Mechanics and Engineering* 190.24–25 (2001), pp. 3247–3270. DOI: 10.1016/S0045-7825(00)00391-1 (cit. on pp. 4, 16, 62).

[22] J. H. Argyris, I. Fried, and D. W. Scharpf. "The TUBA Family of Plate Elements for the Matrix Displacement Method." In: *The Aeronautical Journal* 72.692 (1968), pp. 701–709. DOI: 10.1017/S000192400008489X (cit. on pp. 4, 14).

[23] Philippe G. Ciarlet. *The Finite Element Method for Elliptic Problems*. Vol. 40. Classics in Applied Mathematics. SIAM, 2002 (cit. on pp. 4, 10, 11, 14).

[24] Philippe G. Ciarlet. *Linear and Nonlinear Functional Analysis with Applications*. Other Titles in Applied Mathematics. SIAM, 2013 (cit. on pp. 4, 11, 14).

[25] G. Engel et al. "Continuous/discontinuous finite element approximations of fourth-order elliptic problems in structural and continuum mechanics with applications to thin beams and plates, and strain gradient elasticity." In: *Computer Methods in Applied Mechanics and Engineering* 191.34 (2002), pp. 3669–3750. DOI: 10.1016/S0045-7825(02)00286-4 (cit. on pp. 4, 14, 28, 29, 38, 77, 113).

[26] Susanne C. Brenner and Li-Yeng Sung. "$C^0$ interior penalty methods for fourth order elliptic boundary value problems on polygonal domains." In: *Journal of Scientific Computing* 22.1-3 (2005), pp. 83–118. DOI: 10.1007/s10915-004-4135-7 (cit. on pp. 4, 14, 38, 77, 113).

[27] G.N. Wells and N.T. Dung. "A $C^0$ discontinuous Galerkin formulation for Kirchhoff plates." In: *Computer Methods in Applied Mechanics and Engineering* 196.35–36 (2007), pp. 3370–3380. DOI: 10.1016/j.cma.2007.03.008 (cit. on pp. 4, 14, 28, 29, 38, 77, 83, 113).

[28] Andreas Apostolatos et al. "A Nitsche-type formulation and comparison of the most common domain decomposition methods in isogeometric analysis." In: *International Journal for Numerical Methods in Engineering* 97.7 (2014), pp. 473–504. DOI: 10.1002/nme.4568 (cit. on pp. 5, 13).

[29] Ericka Brivadis et al. "Isogeometric mortar methods." In: *Computer Methods in Applied Mechanics and Engineering* 284 (2015), pp. 292–319. DOI: 10.1016/j.cma.2014.09.012 (cit. on pp. 5, 13).

[30] Austin J. Herrema et al. "Penalty coupling of non-matching isogeometric Kirchhoff–Love shell patches with application to composite wind turbine blades." In: *Computer Methods in Applied Mechanics and Engineering* 346 (2019), pp. 810–840. DOI: 10.1016/j.cma.2018.08.038 (cit. on pp. 5, 13).

[31] L. Coradello, G. Loli, and A. Buffa. "A projected super-penalty method for the $C^1$-coupling of multi-patch isogeometric Kirchhoff plates." In: *Computational Mechanics* 67 (2021), pp. 1133–1153. DOI: 10.1007/s00466-021-01983-w (cit. on pp. 5, 13, 30).

[32] M. Chasapi, P. Antolín, and A. Buffa. "Fast parametric analysis of trimmed multi-patch isogeometric Kirchhoff–Love shells using a local reduced basis method." In: *Engineering with Computers* 40 (2024), pp. 3623–3650. DOI: 10.1007/s00366-024-01980-6 (cit. on pp. 5, 13, 30).

[33] Michael Hintermüller, Kazufumi Ito, and Karl Kunisch. "The Primal-Dual Active Set Strategy as a Semismooth Newton Method." In: *SIAM Journal on Optimization* 13.3 (2003), pp. 865–888. DOI: 10.1137/S1052623401383558 (cit. on pp. 5, 32).

[34] Nathan M. Newmark. "A Method of Computation for Structural Dynamics." In: *Journal of the Engineering Mechanics Division* 85.3 (1959), pp. 67–94. DOI: 10.1061/JMCEA3.0000098 (cit. on pp. 5, 16, 115).

[35] William L. Oberkampf and Christopher J. Roy. *Verification, Validation, and Uncertainty Quantification in Scientific Computing*. 2nd ed. Cambridge University Press, 2025 (cit. on pp. 5, 15, 33, 85, 125).

[36] The American Society of Mechanical Engineers. *Standard for Verification and Validation in Computational Solid Mechanics*. ASME V&V 10–2019. New York: The American Society of Mechanical Engineers, 2019 (cit. on pp. 5, 15, 33, 85, 125).

[37] Anders Logg, Kent-Andre Mardal, and Garth N. Wells, eds. *Automated Solution of Differential Equations by the Finite Element Method: The FEniCS Book*. Vol. 84. Lecture Notes in Computational Science and Engineering. Berlin, Heidelberg: Springer, 2012. DOI: 10.1007/978-3-642-23099-8 (cit. on pp. 5, 79, 146).

[38] J.S. Hale et al. "Simple and extensible plate and shell finite element models through automatic code generation tools." In: *Computers & Structures* 209 (2018), pp. 163–181. DOI: 10.1016/j.compstruc.2018.08.001 (cit. on pp. 5, 77, 113).

[39] Peter Benner, Serkan Gugercin, and Karen Willcox. "A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems." In: *SIAM Review* 57.4 (2015), pp. 483–531. DOI: 10.1137/130932715 (cit. on pp. 6, 10, 18, 42, 127).

[40] Alfio Quarteroni, Andrea Manzoni, and Federico Negri. *Reduced Basis Methods for Partial Differential Equations: An Introduction*. Vol. 92. UNITEXT. Cham: Springer, 2016 (cit. on pp. 6, 10, 17, 18, 42, 89, 127).

[41] J.S. Hesthaven, G. Rozza, and B. Stamm. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. Springer, 2016 (cit. on pp. 6, 17, 18, 42, 89, 127).

[42] Peter Benner et al., eds. *Model Order Reduction. Volume 2: Snapshot-Based Methods and Algorithms*. Berlin, Boston: De Gruyter, 2021. DOI: 10.1515/9783110671490 (cit. on p. 6).

[43] G. Rozza et al. *Real Time Reduced Order Computational Mechanics: Parametric PDEs Worked Out Problems*. SISSA Springer Series. Cham: Springer Nature Switzerland, 2024 (cit. on pp. 6, 18, 42, 146).

[44] Fabrizio Gelsomino and Gianluigi Rozza. "Comparison and Combination of Reduced-Order Modelling Techniques in 3D Parametrized Heat Transfer Problems." In: *Mathematical and Computer Modelling of Dynamical Systems* 17.4 (2011), pp. 371–394. DOI: 10.1080/13873954.2011.547672 (cit. on pp. 6, 87).

[45] Nirav Vasant Shah et al. "Finite element based model order reduction for parametrized one-way coupled steady state linear thermo-mechanical problems." In: *Finite Elements in Analysis and Design* 212 (2022), p. 103837. DOI: 10.1016/j.finel.2022.103837 (cit. on pp. 6, 87).

[46] Nicola Demo, Marco Tezzele, and Gianluigi Rozza. "A non-intrusive approach for the reconstruction of POD modal coefficients through active subspaces." In: *Comptes Rendus. Mécanique* 347.11 (2019), pp. 873–881. DOI: 10.1016/j.crme.2019.11.012 (cit. on pp. 6, 18, 42, 44, 45).

[47] J.S. Hesthaven and S. Ubbiali. "Non-intrusive reduced order modeling of nonlinear problems using neural networks." In: *Journal of Computational Physics* 363 (2018), pp. 55–78. DOI: 10.1016/j.jcp.2018.02.037 (cit. on pp. 6, 18, 42, 44, 46, 92, 132).

[48] Mengwu Guo and Jan S. Hesthaven. "Reduced order modeling for nonlinear structural analysis using Gaussian process regression." In: *Computer Methods in Applied Mechanics and Engineering* 341 (2018), pp. 807–826. DOI: 10.1016/j.cma.2018.07.017 (cit. on pp. 6, 42, 47, 91).

[49] Stefania Fresca and Andrea Manzoni. "POD-DL-ROM: Enhancing Deep Learning-Based Reduced Order Models for Nonlinear Parametrized PDEs by Proper Orthogonal Decomposition." In: *Computer Methods in Applied Mechanics and Engineering* 388 (2022), p. 114181. DOI: 10.1016/j.cma.2021.114181 (cit. on pp. 6, 18, 92, 133).

[50] Stefania Fresca, Luca Dedè, and Andrea Manzoni. "A Comprehensive Deep Learning-Based Approach to Reduced Order Modeling of Nonlinear Time-Dependent Parametrized PDEs." In: *Journal of Scientific Computing* 87 (2021), p. 61. DOI: 10.1007/s10915-021-01462-7 (cit. on pp. 6, 18, 92).

[51] Arthur W. Leissa. *Vibration of Plates*. NASA Special Publication SP-160. Washington, DC: National Aeronautics and Space Administration, 1969 (cit. on pp. 16, 116).

[52] T. K. Caughey and M. E. J. O'Kelly. "Classical Normal Modes in Damped Linear Dynamic Systems." In: *Journal of Applied Mechanics* 32.3 (1965), pp. 583–588. DOI: 10.1115/1.3627262 (cit. on pp. 16, 108).

[53] Lawrence Sirovich. "Turbulence and the Dynamics of Coherent Structures. Part I: Coherent Structures." In: *Quarterly of Applied Mathematics* 45.3 (1987), pp. 561–571 (cit. on pp. 17, 43, 89, 127).

[54] Shachar Kaufman et al. "Leakage in Data Mining: Formulation, Detection, and Avoidance." In: *ACM Transactions on Knowledge Discovery from Data* 6.4, 15 (2012), pp. 1–21. DOI: 10.1145/2382577.2382579 (cit. on pp. 18, 48, 89, 127).

[55] David R. Roberts et al. "Cross-Validation Strategies for Data with Temporal, Spatial, Hierarchical, or Phylogenetic Structure." In: *Ecography* 40.8 (2017), pp. 913–929. DOI: 10.1111/ecog.02881 (cit. on pp. 18, 127).

[56] Andrea Manzoni, Alfio Quarteroni, and Gianluigi Rozza. "Computational Reduction for Parametrized PDEs: Strategies and Applications." In: *Milan Journal of Mathematics* 80.2 (2012), pp. 283–309. DOI: 10.1007/s00032-012-0182-y (cit. on p. 18).

[57] Maxime Barrault et al. "An empirical interpolation method: application to efficient reduced-basis discretization of partial differential equations." In: *Comptes Rendus Mathematique* 339.9 (2004), pp. 667–672. DOI: 10.1016/j.crma.2004.08.006 (cit. on p. 18).

[58] Saifon Chaturantabut and Danny C. Sorensen. "Nonlinear model reduction via discrete empirical interpolation." In: *SIAM Journal on Scientific Computing* 32.5 (2010), pp. 2737–2764. DOI: 10.1137/090766498 (cit. on p. 18).

[59] Michael A. Grepl et al. "Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations." In: *ESAIM: Mathematical Modelling and Numerical Analysis* 41.3 (2007), pp. 575–605. DOI: 10.1051/m2an:2007031 (cit. on p. 18).

[60] Federico Pichi et al. "An artificial neural network approach to bifurcating phenomena in computational fluid dynamics." In: *Computers & Fluids* 254 (2023), p. 105813. DOI: 10.1016/j.compfluid.2023.105813 (cit. on pp. 18, 92).

[61] Federico Pichi, Beatriz Moya, and Jan S. Hesthaven. "A Graph Convolutional Autoencoder Approach to Model Order Reduction for Parametrized PDEs." In: *Journal of Computational Physics* 501 (2024), p. 112762. DOI: 10.1016/j.jcp.2024.112762 (cit. on p. 18).

[62] Philippe G. Ciarlet. *Mathematical Elasticity, Volume II: Theory of Plates*. Vol. 27. Studies in Mathematics and Its Applications. Amsterdam: North-Holland, 1997 (cit. on p. 22).

[63] Y. H. Wang, L. G. Tham, and Y. K. Cheung. "Beams and Plates on Elastic Foundations: A Review." In: *Progress in Structural Engineering and Materials* 7.4 (2005), pp. 174–182. DOI: 10.1002/pse.202 (cit. on p. 23).

[64] Piero Villaggio. "A Free Boundary Value Problem in Plate Theory." In: *Journal of Applied Mechanics* 50.2 (1983), pp. 297–302. DOI: 10.1115/1.3167035 (cit. on p. 23).

[65] Ivo Babuška. "The Finite Element Method with Penalty." In: *Mathematics of Computation* 27.122 (1973), pp. 221–228. DOI: 10.1090/S0025-5718-1973-0351118-5 (cit. on p. 24).

[66] Douglas N. Arnold. "An Interior Penalty Finite Element Method with Discontinuous Elements." In: *SIAM Journal on Numerical Analysis* 19.4 (1982), pp. 742–760. DOI: 10.1137/0719052 (cit. on p. 29).

[67] Peter Hansbo and Mats G. Larson. "A Discontinuous Galerkin Method for the Plate Equation." In: *Calcolo* 39.1 (2002), pp. 41–59. DOI: 10.1007/s100920200001 (cit. on p. 29).

[68] Khosro Shahbazi. "An Explicit Expression for the Penalty Parameter of the Interior Penalty Method." In: *Journal of Computational Physics* 205.2 (2005), pp. 401–407. DOI: 10.1016/j.jcp.2004.11.017 (cit. on p. 30).

[69] Andréa R. D. Silva, Ricardo A. M. Silveira, and Paulo B. Gonçalves. "Numerical Methods for Analysis of Plates on Tensionless Elastic Foundations." In: *International Journal of Solids and Structures* 38.10–13 (2001), pp. 2083–2100. DOI: 10.1016/S0020-7683(00)00154-2 (cit. on p. 31).

[70] Noboru Kikuchi and J. Tinsley Oden. *Contact Problems in Elasticity: A Study of Variational Inequalities and Finite Element Methods*. Vol. 8. Studies in Applied and Numerical Mathematics. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1988 (cit. on p. 31).

[71] Alfio Quarteroni. *Numerical Models for Differential Problems*. 3rd ed. Vol. 16. MS&A: Modeling, Simulation and Applications. Springer, 2017 (cit. on p. 32).

[72] Peter Wriggers. *Computational Contact Mechanics*. 2nd ed. Berlin, Heidelberg: Springer, 2006 (cit. on p. 32).

[73] Ivo Babuška and J. Tinsley Oden. "Verification and Validation in Computational Engineering and Science: Basic Concepts." In: *Computer Methods in Applied Mechanics and Engineering* 193.36–38 (2004), pp. 4057–4066. DOI: 10.1016/j.cma.2004.03.002 (cit. on p. 33).

[74] Christophe Prud'homme et al. "Reliable Real-Time Solution of Parametrized Partial Differential Equations: Reduced-Basis Output Bound Methods." In: *Journal of Fluids Engineering* 124.1 (2002), pp. 70–80. DOI: 10.1115/1.1448332 (cit. on p. 42).

[75] Roberto Milani, Alfio Quarteroni, and Gianluigi Rozza. "Reduced Basis Method for Linear Elasticity Problems with Many Parameters." In: *Computer Methods in Applied Mechanics and Engineering* 197.51–52 (2008), pp. 4812–4829. DOI: 10.1016/j.cma.2008.07.002 (cit. on p. 42).

[76] Dinh Bao Phuong Huynh, Federico Pichi, and Gianluigi Rozza. "Reduced Basis Approximation and A Posteriori Error Estimation: Applications to Elasticity Problems in Several Parametric Settings." In: *Numerical Methods for PDEs: State of the Art Techniques*. Ed. by Daniele Antonio Di Pietro, Alexandre Ern, and Luca Formaggia. Vol. 15. SEMA SIMAI Springer Series. Cham: Springer International Publishing, 2018, pp. 203–247. DOI: 10.1007/978-3-319-94676-4_8 (cit. on p. 42).

[77] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press, 2006 (cit. on pp. 42, 47, 91, 132).

[78] M. D. McKay, Richard J. Beckman, and William J. Conover. "A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code." In: *Technometrics* 21.2 (1979), pp. 239–245. DOI: 10.1080/00401706.1979.10489755 (cit. on p. 42).

[79] Jerome Sacks et al. "Design and Analysis of Computer Experiments." In: *Statistical Science* 4.4 (1989), pp. 409–423. DOI: 10.1214/ss/1177012413 (cit. on p. 42).

[80] Maciej Balajewicz, David Amsallem, and Charbel Farhat. "Projection-Based Model Reduction for Contact Problems." In: *International Journal for Numerical Methods in Engineering* 106.8 (2016), pp. 644–663. DOI: 10.1002/nme.5135 (cit. on p. 44).

[81] D. Xiao et al. "Non-Intrusive Reduced-Order Modelling of the Navier–Stokes Equations Based on RBF Interpolation." In: *International Journal for Numerical Methods in Fluids* 79.11 (2015), pp. 580–595. DOI: 10.1002/fld.4066 (cit. on pp. 46, 91, 132).

[82] Holger Wendland. *Scattered Data Approximation*. Vol. 17. Cambridge Monographs on Applied and Computational Mathematics. Cambridge: Cambridge University Press, 2004 (cit. on pp. 46, 91, 132).

[83] Mark de Berg et al. *Computational Geometry: Algorithms and Applications*. 3rd ed. Berlin, Heidelberg: Springer, 2008 (cit. on p. 46).

[84] C. Bradford Barber, David P. Dobkin, and Hannu Huhdanpaa. "The Quickhull Algorithm for Convex Hulls." In: *ACM Transactions on Mathematical Software* 22.4 (1996), pp. 469–483. DOI: 10.1145/235815.235821 (cit. on p. 46).

[85] Francisco Armero and Juan C. Simo. "A New Unconditionally Stable Fractional Step Method for Non-Linear Coupled Thermomechanical Problems." In: *International Journal for Numerical Methods in Engineering* 35.4 (1992), pp. 737–766. DOI: 10.1002/nme.1620350408 (cit. on p. 62).

[86] European Commission, Joint Research Centre. *PVGIS: Photovoltaic Geographical Information System*. Accessed: 2026-04. 2024. URL: https://joint-research-centre.ec.europa.eu/photovoltaic-geographical-information-system-pvgis_en (cit. on p. 65).

[87] ENEA. *Solar Radiation Tables for Italy*. Accessed: 2026-04. 2024. URL: https://www.solaritaly.enea.it/TabelleRad/TabelleRadEn.php (cit. on p. 65).

[88] Theodore R. Tauchert. "Thermally Induced Flexure, Buckling, and Vibration of Plates." In: *Applied Mechanics Reviews* 44.8 (1991), pp. 347–360. DOI: 10.1115/1.3119508 (cit. on p. 66).

[89] J. F. Stelzer and R. Welzel. "Experiences in Non-Linear Analysis of Temperature Fields with Finite Elements." In: *International Journal for Numerical Methods in Engineering* 24.1 (1987), pp. 59–73. DOI: 10.1002/nme.1620240105 (cit. on p. 75).

[90] Xiangmin Jiao and Michael T. Heath. "Common-Refinement-Based Data Transfer between Non-Matching Meshes in Multiphysics Simulations." In: *International Journal for Numerical Methods in Engineering* 61.14 (2004), pp. 2402–2427. DOI: 10.1002/nme.1147 (cit. on p. 75).

[91] Alexander M. Steenhoek, Daniel J. Rixen, and Philippe Nachtergaele. "Static Correction in Model Order Reduction Techniques for Multiphysical Problems." In: *Structural Dynamics, Volume 3*. Ed. by Tom Proulx. Conference Proceedings of the Society for Experimental Mechanics Series. New York, NY: Springer, 2011, pp. 647–662. DOI: 10.1007/978-1-4419-9834-7_57 (cit. on p. 87).

[92] Timo Reis and Tatjana Stykel. "A Survey on Model Reduction of Coupled Systems." In: *Model Order Reduction: Theory, Research Aspects and Applications*. Ed. by Wil H. A. Schilders, Henk A. van der Vorst, and Joost Rommes. Vol. 13. Mathematics in Industry. Berlin, Heidelberg: Springer, 2008, pp. 133–155. DOI: 10.1007/978-3-540-78841-6_7 (cit. on p. 91).

[93] Giulio Ortali, Nicola Demo, and Gianluigi Rozza. "A Gaussian Process Regression Approach within a Data-Driven POD Framework for Engineering Problems in Fluid Dynamics." In: *Mathematics in Engineering* 4.3 (2022), pp. 1–16. DOI: 10.3934/mine.2022021 (cit. on p. 91).

[94] Raymond D. Mindlin. "Influence of Rotatory Inertia and Shear on Flexural Motions of Isotropic, Elastic Plates." In: *Journal of Applied Mechanics* 18.1 (1951), pp. 31–38. DOI: 10.1115/1.4010217 (cit. on p. 102).

[95] Michel Géradin and Daniel J. Rixen. *Mechanical Vibrations: Theory and Application to Structural Dynamics*. 3rd ed. Chichester: John Wiley & Sons, 2015 (cit. on pp. 107, 108, 115, 116).

[96] J. L. Lions and E. Magenes. "Evolution Equations of the Second Order in $t$ and of Schroedinger Type." In: *Non-Homogeneous Boundary Value Problems and Applications*. Vol. 183. Die Grundlehren der mathematischen Wissenschaften. Berlin, Heidelberg: Springer, 1973, pp. 246–278. DOI: 10.1007/978-3-642-65393-3_5 (cit. on p. 107).

[97] Thomas J. R. Hughes. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Mineola, NY: Dover Publications, 2000 (cit. on p. 114).

[98] E. Hinton, T. Rock, and O. C. Zienkiewicz. "A Note on Mass Lumping and Related Processes in the Finite Element Method." In: *Earthquake Engineering & Structural Dynamics* 4.3 (1976), pp. 245–249. DOI: 10.1002/eqe.4290040305 (cit. on p. 114).

[99] Klaus-Jürgen Bathe and Edward L. Wilson. "Stability and Accuracy Analysis of Direct Integration Methods." In: *Earthquake Engineering & Structural Dynamics* 1.3 (1972), pp. 283–291. DOI: 10.1002/eqe.4290010308 (cit. on p. 115).

[100] Klaus-Jürgen Bathe and Edward L. Wilson. "Solution Methods for Eigenvalue Problems in Structural Mechanics." In: *International Journal for Numerical Methods in Engineering* 6.2 (1973), pp. 213–226. DOI: 10.1002/nme.1620060207 (cit. on p. 116).

[101] Steen Krenk. "Energy Conservation in Newmark Based Time Integration Algorithms." In: *Computer Methods in Applied Mechanics and Engineering* 195.44–47 (2006), pp. 6110–6124. DOI: 10.1016/j.cma.2005.12.001 (cit. on p. 123).

[102] Jan S. Hesthaven, Cecilia Pagliantini, and Gianluigi Rozza. "Reduced Basis Methods for Time-Dependent Problems." In: *Acta Numerica* 31 (2022), pp. 265–345. DOI: 10.1017/S0962492922000058 (cit. on p. 127).

[103] Mengwu Guo and Jan S. Hesthaven. "Data-Driven Reduced Order Modeling for Time-Dependent Problems." In: *Computer Methods in Applied Mechanics and Engineering* 345 (2019), pp. 75–99. DOI: 10.1016/j.cma.2018.10.029 (cit. on pp. 127, 132).

[104] Christophe Audouze, Florian De Vuyst, and Prasanth B. Nair. "Nonintrusive Reduced-Order Modeling of Parametrized Time-Dependent Partial Differential Equations." In: *Numerical Methods for Partial Differential Equations* 29.5 (2013), pp. 1587–1628. DOI: 10.1002/num.21768 (cit. on pp. 127, 132).

[105] Karl Kunisch and Stefan Volkwein. "Galerkin Proper Orthogonal Decomposition Methods for Parabolic Problems." In: *Numerische Mathematik* 90 (2001), pp. 117–148. DOI: 10.1007/s002110100282 (cit. on p. 130).

[106] Gaëtan Kerschen et al. "The Method of Proper Orthogonal Decomposition for Dynamical Characterization and Order Reduction of Mechanical Systems: An Overview." In: *Nonlinear Dynamics* 41.1–3 (2005), pp. 147–169. DOI: 10.1007/s11071-005-2803-2 (cit. on p. 130).

[107] Qian Wang, Jan S. Hesthaven, and Deep Ray. "Non-Intrusive Reduced Order Modeling of Unsteady Flows Using Artificial Neural Networks with Application to a Combustion Problem." In: *Journal of Computational Physics* 384 (2019), pp. 289–307. DOI: 10.1016/j.jcp.2019.01.031 (cit. on p. 132).

[108] Geir Kjetil Sandve et al. "Ten Simple Rules for Reproducible Computational Research." In: *PLoS Computational Biology* 9.10 (2013), e1003285. DOI: 10.1371/journal.pcbi.1003285 (cit. on p. 146).

[109] Greg Wilson et al. "Good Enough Practices in Scientific Computing." In: *PLoS Computational Biology* 13.6 (2017), e1005510. DOI: 10.1371/journal.pcbi.1005510 (cit. on p. 146).

[110] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009 (cit. on p. 146).

[111] Martin S. Alnæs. "UFL: A Finite Element Form Language." In: *Automated Solution of Differential Equations by the Finite Element Method*. Ed. by Anders Logg, Kent-Andre Mardal, and Garth N. Wells. Vol. 84. Lecture Notes in Computational Science and Engineering. Berlin, Heidelberg: Springer, 2012, pp. 303–338. DOI: 10.1007/978-3-642-23099-8_17 (cit. on p. 146).

[112] Charles R. Harris et al. "Array Programming with NumPy." In: *Nature* 585.7825 (2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2 (cit. on p. 146).

[113] Pauli Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python." In: *Nature Methods* 17.3 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2 (cit. on p. 146).

[114] Fabian Pedregosa et al. "Scikit-learn: Machine Learning in Python." In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830 (cit. on p. 146).

[115] Adam Paszke et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library.” In: *Advances in Neural Information Processing Systems*. Vol. 32. 2019 (cit. on p. 146).

[116] Wes McKinney. “Data Structures for Statistical Computing in Python.” In: *Proceedings of the 9th Python in Science Conference*. 2010, pp. 56–61. DOI: 10.25080/Majora-92bf1922-00a (cit. on p. 146).

[117] John D. Hunter. “Matplotlib: A 2D Graphics Environment.” In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. DOI: 10.1109/MCSE.2007.55 (cit. on p. 146).

[118] RBniCS Project. *RBniCS: Reduced Order Modelling in FEniCS*. 2026. URL: https://www.rbnicsproject.org/ (visited on 05/16/2026) (cit. on p. 146).

[119] Thomas Kluyver et al. “Jupyter Notebooks: A Publishing Format for Reproducible Computational Workflows.” In: *Positioning and Power in Academic Publishing: Players, Agents and Agendas*. Ed. by Fernando Loizides and Birgit Schmidt. IOS Press, 2016, pp. 87–90. DOI: 10.3233/978-1-61499-649-1-87 (cit. on p. 146).