



Von Mises stress concentrating at a hole in a loaded plate — linear elasticity, verified in DOLFINx.

Verified Finite Elements with DOLFIN_x

*Solving PDEs with the FEniCSx Stack, Checked by
Manufactured Solutions*

Alexios Zervos



ODIN PRESS



Verified Finite Elements with DOLFINx

Copyright © 2026 by **Alexios Zervos**.

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopying, recording, or otherwise—without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

Paperback version, First edition: 2026.

ISBN: forthcoming (Bowker allocation pending)

Trim Size: 6 × 9 inch. Paper color: White.

Source-code copyright. All Python source code listings reproduced in this volume—including those in the Appendix, the chapter bodies, and any auxiliary file distributed alongside this book—are the copyrighted intellectual property of **Alexios Zervos**. No reproduction, redistribution, adaptation, translation into another programming language, or incorporation into any derivative work, whether academic, commercial, or otherwise, is permitted without the prior written permission of the copyright holder. Single-machine execution by the individual purchaser of this volume for the sole purpose of reproducing a figure or table from this book is granted as a limited, non-transferable, non-sublicensable use and does not constitute redistribution.

Published by Odin Press.

<https://odinpress.org>

This book was typeset in L^AT_EX by Odin Press from a manuscript submitted by the author.

Printed and Bound in the UNITED STATES OF AMERICA.

9 8 7 6 5 4 3 2 1 0

Preface

A finite-element simulation you cannot verify is an opinion, not a result. I have believed this for forty years, and nothing in the migration of the finite-element community onto DOLFINx has given me cause to soften it. A new solver, a cleaner library, an assembly kernel far faster than the one before — these are welcome, and this book is built on them. But a new tool is exactly when a discipline’s verification habits lapse: the tool is exciting and the checking is not, and a smooth-looking stress field asks no questions of the person who produced it. This book exists so the move to DOLFINx does not cost us the one habit that separates a computed result from a guess: forcing a solver to reproduce an answer we already know before we trust one we do not.

The gap this book fills is not in the methods: the finite-element method is old, its theory settled, DOLFINx implements it well. The gap is that the methods are almost always taught as if producing a field were the end of the work, when it is only the beginning. A textbook shows the weak form, assembles the matrix, and plots the solution; it rarely measures whether that solution converges at the order the theory promised, and almost never for every method it presents. The commercial packages are worse, not because their elements are wrong but because you cannot audit them — I have watched two disagree by fifteen percent on the same thermal-stress problem, neither willing to show me its quadrature. An unauditible result is untrustworthy, however professional the picture. What has been missing is a book that solves real problems with a modern open stack and verifies every one of them, in the open, by a method the reader can run himself.

So that is what this book does, the same way throughout: every method is verified before it is trusted. Where a problem has a closed-form solution — a Lamé cylinder, a Timoshenko bimetallic strip — the computed answer is matched against it. Where it does not, and most real geometries do not, the code is verified by the method of manufactured solutions: cook a

solution, derive the forcing it satisfies, and measure whether the error falls at the a-priori rate the theory promises. The convergence table goes on the page: I do not report "good agreement," I report the slope and let it make the argument. Where neither a closed form nor a manufactured solution settles the matter, the physics does — a reaction-force balance that must sum to zero, a peak stress that must stop moving as the mesh refines, the same answer on one processor and on eight. The book is candid, too, about where the method quietly fails: the volumetric locking that stiffens a low-order element to nonsense without warning, the checkerboard pressure of an inf-sup pair chosen by habit, the geometry error that hides in a functional rather than a solution slope. Every figure is produced by one self-contained file against a pinned library version; a figure you cannot regenerate is an anecdote, and there are none here.

The book is arranged as an ascent. The opening chapters build the discipline: the weak form and the assembly it drives, the method of manufactured solutions and the order it measures, the boundary conditions that must be imposed exactly, the unstructured meshes that come from a real CAD pipeline, and the elements whose interpolation error sets every rate that follows. The middle chapters put that discipline to work on the physics: linear elasticity and the locking that manufactured solutions expose, time-dependent heat and the two rates a transient carries, the Stokes saddle point and the inf-sup condition that governs it, and the nonlinear problems whose Newton iteration converges quadratically only when the Jacobian is exact. The computational chapters make the whole thing scale: the multigrid solver that cuts the cost without touching the answer, the parallelism whose rank-independence is itself a verification check, and the reproducible workflow that turns a claim into a test anyone can rerun. The book then spends all of it at once on a single real part — an irregular steel wrench, meshed at its fillets, heated until its steel softens and loaded until it yields, verified by closed form where one exists and by convergence and conservation where none does, and reported honestly down to the safety factor below one that says the red-hot part cannot take the load.

I came to DOLFINx late and with one instinct: never trust a solver you cannot force to reproduce an answer you already know. This book is that instinct worked out in full, on a tool worth trusting once you have made it prove the order first.

Αλέξιος Ζερόβος

Contents

Preface	ix
List of Figures	xvii
List of Tables	xxv
I Foundations of Verified Finite Elements	1
1 The Weak Form and Your First Solve: Poisson from Space to Solution	3
1.1 The Poisson Problem and Its Weak Form	4
1.2 Function Spaces and the Discrete Galerkin Problem . . .	8
1.3 Dirichlet Boundary Conditions	13
1.4 Assembling the Sparse Linear System from UFL	15
1.5 Solving with PETSc: SPD Systems and the Conjugate Gradient Method	19
1.6 Verification by Manufactured Solutions: Measuring the Convergence Order	22
2 Verification as a Discipline: Manufactured Solutions and the Order $p+1$	31
2.1 Verification versus Validation	32
2.2 The Method of Manufactured Solutions	34
2.3 A-priori Error Estimates and the Expected Orders . . .	38
2.4 Measuring the Observed Order of Accuracy	41
2.5 The Reusable Verification Toolkit	46
2.6 What Verification Does Not Catch	48
3 Boundary Conditions, Done Right: Dirichlet, Neumann, and Robin, Each Verified	53

3.1	Essential and Natural: The Two Kinds of Boundary Condition	54
3.2	Dirichlet Conditions on the Trial Space	57
3.3	Neumann Conditions in the Weak Form	60
3.4	Robin Conditions on the Boundary	64
3.5	Multiple and Mixed Conditions on One Boundary	68
3.6	Verifying Each Condition by Manufactured Solutions	70
4	Meshes and Geometry: From Built-In Grids to gmsh Import and h-Refinement	75
4.1	Built-In Meshes and the Mesh as a Cell Complex	76
4.2	Cells, the Reference Cell, and the Geometry Map	80
4.3	Unstructured Meshes with gmsh	84
4.4	Importing a gmsh Mesh into DOLFINx	87
4.5	Higher-Order and Curved Geometry	90
4.6	h-Refinement Convergence, Verified	95
II	The Core Problem Types	99
5	Function Spaces and Elements: Lagrange, Higher Order, and the Basix Zoo	101
5.1	The Finite Element: Cell, Space, and Degrees of Freedom	102
5.2	Lagrange Elements of Arbitrary Degree	104
5.3	Interpolation and Projection into a Space	109
5.4	The Basix Element Zoo	112
5.5	p-Refinement: Raising the Degree	114
5.6	Verifying Spaces: Interpolation Error and p-Convergence	116
6	Vector Problems: Linear Elasticity, the Patch Test, and Volumetric Locking	121
6.1	The Equations of Linear Elasticity	122
6.2	The Weak Form and Vector Function Spaces	124
6.3	The Patch Test	129
6.4	Verification by Manufactured Solutions	132
6.5	Volumetric Locking Near Incompressibility	135
6.6	Locking Remedies, Verified	140
7	Time-Dependent Problems: The Heat Equation in Space and Time	145

7.1	The Heat Equation and Its Spatial Weak Form	146
7.2	Time Discretization: The Theta-Method	149
7.3	The Fully Discrete Scheme: Assemble Once, Reuse	152
7.4	Stability and Energy Decay	155
7.5	Verification: Spatial and Temporal Order	159
7.6	A Worked Transient Problem	162
8	Mixed and Saddle-Point Problems: Stokes, the inf-sup Condition, and Taylor–Hood	167
8.1	The Stokes Problem and Its Saddle-Point Structure	168
8.2	Mixed Velocity-Pressure Spaces	171
8.3	The inf-sup (LBB) Condition	172
8.4	Stable and Unstable Pairs	175
8.5	Solving the Saddle-Point System	179
8.6	Verification and the Checkerboard Failure	181
9	Nonlinear Problems: Newton’s Method and the Automatic Jacobian	187
9.1	Nonlinear Weak Forms	188
9.2	Newton’s Method for the Discrete Residual	189
9.3	The Automatic Jacobian	192
9.4	Quadratic Convergence and Its Verification	194
9.5	PETSc SNES and Globalization	196
9.6	A Worked Nonlinear Problem	198
III	Solvers, Parallelism, and Reproducibility	205
10	Solvers and Preconditioning with PETSc: From Direct Factorization to Field-Split	207
10.1	Direct Solvers and Their 3D Scaling	208
10.2	Krylov Methods: CG and GMRES	211
10.3	Preconditioning and the Condition Number	212
10.4	Algebraic Multigrid as an Optimal Preconditioner	215
10.5	Field-Split for Block Systems	216
10.6	Verifying the Solver: Same Answer, Better Scaling	218
11	Parallelism and Performance: MPI by Default and Rank-Independence as a Check	223
11.1	The Distributed Mesh: Partitioning and Ghost Cells	224

11.2	MPI by Default: The SPMD Model	226
11.3	Parallel Assembly and Ghost Updates	228
11.4	Parallel Solvers and Their Communication	230
11.5	Rank-Independence as a Verification Check	232
11.6	Strong and Weak Scaling	234
12	Reproducible Workflows: I/O, Checkpointing, and the Verification Dashboard	241
12.1	The Reproducibility Contract	242
12.2	I/O: Writing Fields and Meshes	245
12.3	Checkpointing and Restart	248
12.4	The Verification Dashboard	250
12.5	Containers and Pinned Environments	254
12.6	A Reproducible Verified Run	256
IV	Capstone: A Verified Engineering Analysis	261
13	The Steel Wrench: Geometry, Mesh, and the One-Way Thermo-Mechanical Solver	263
13.1	The Steel Wrench: Geometry and the gmsh Model	264
13.2	Meshing the Wrench	266
13.3	Linear Thermoelasticity: The One-Way Coupling	268
13.4	Real Steel Constants and the Load Case	271
13.5	Verification on Companion Shapes	274
13.6	The One-Way Solver Assembled and Checked	278
14	The Real Part, Verified: MMS on the Wrench and Temperature-Dependent Steel	283
14.1	Manufactured Solutions on the Wrench Mesh	284
14.2	Quantity-of-Interest Convergence to a Fine Reference	286
14.3	Physical Sanity: Reaction-Force Balance and Symmetry	288
14.4	Temperature-Dependent Steel Properties	290
14.5	The Mild Nonlinearity: Newton	293
14.6	Verification Against the Constant-Property Baseline	295
15	At Scale and the Full Report: Transient Heating, Combined Load, and the Dashboard	301
15.1	Transient Heating: The Evolving Thermal Stress	302
15.2	The Combined Load Case	305

15.3	At Scale: Parallelism and Rank-Independence on the Wrench	308
15.4	The Verification Dashboard: The Full Report	310
15.5	The Verification Ledger: The Provenance of Every Number	313
15.6	The Graduation Piece	315
A Source Code		321
A.1	Chapter 1 — The Weak Form and Your First Solve: Poison from Space to Solution	321
A.2	Chapter 2 — Verification as a Discipline: Manufactured Solutions and the Order $p+1$	329
A.3	Chapter 3 — Boundary Conditions, Done Right: Dirichlet, Neumann, and Robin, Each Verified	335
A.4	Chapter 4 — Meshes and Geometry: From Built-In Grids to gmsh Import and h-Refinement	342
A.5	Chapter 5 — Function Spaces and Elements: Lagrange, Higher Order, and the Basix Zoo	354
A.6	Chapter 6 — Vector Problems: Linear Elasticity, the Patch Test, and Volumetric Locking	361
A.7	Chapter 7 — Time-Dependent Problems: The Heat Equation in Space and Time	371
A.8	Chapter 8 — Mixed and Saddle-Point Problems: Stokes, the inf-sup Condition, and Taylor–Hood	383
A.9	Chapter 9 — Nonlinear Problems: Newton’s Method and the Automatic Jacobian	392
A.10	Chapter 10 — Solvers and Preconditioning with PETSc: From Direct Factorization to Field-Split	400
A.11	Chapter 11 — Parallelism and Performance: MPI by Default and Rank-Independence as a Check	408
A.12	Chapter 12 — Reproducible Workflows: I/O, Checkpointing, and the Verification Dashboard	415
A.13	Chapter 13 — The Steel Wrench: Geometry, Mesh, and the One-Way Thermo-Mechanical Solver	422
A.14	Chapter 14 — The Real Part, Verified: MMS on the Wrench and Temperature-Dependent Steel	440
A.15	Chapter 15 — At Scale and the Full Report: Transient Heating, Combined Load, and the Dashboard	452
References		469

Contents

Index	481
About the Author	485

Chapter 1

The Weak Form and Your First Solve: Poisson from Space to Solution

A finite-element program that produces a smooth, colourful picture has told you nothing you can trust. This book is built on the opposite discipline: a solver earns belief only when it reproduces an answer you already knew, at the rate the theory promised, and it earns that belief once per method and never permanently. The engine that supplies the known answers is the method of manufactured solutions, and the yardstick is the measured order of convergence — the slope of error against mesh size, read off a table and reported as a number. Every chapter that follows solves a partial differential equation with DOLFINx and the FEniCSx stack [1], and every chapter verifies the solve before it believes it.

The Poisson problem is where the whole apparatus assembles for the first time. It is the simplest elliptic equation with a genuine boundary-value structure, it is symmetric and positive definite so its linear algebra is the friendliest we will meet, and it has manufactured solutions of every smoothness we could want. That makes it the right place to lay the six pieces the rest of the book reuses: the weak form the finite-element method actually solves, the finite-dimensional space that discretizes it, the Dirichlet condition that pins the solution to its data, the sparse system that assembly produces, the PETSc Krylov solver that inverts it [2], and the manufactured-solution check that certifies the result. This chapter builds

each piece, proves the two theorems that make the target well posed and its discretization quasi-optimal, and closes with the convergence study that turns the pipeline from an assertion into a measurement. When the L^2 error falls by four and the H^1 error by two each time the mesh is halved, the solver has passed its first verification — and only then is it a result rather than an opinion.

1.1 The Poisson Problem and Its Weak Form

Let $\Omega \subset \mathbb{R}^2$ be a bounded domain with Lipschitz boundary $\partial\Omega$, and let f be a given source and g given boundary data. The Poisson problem asks for a scalar field whose negative Laplacian matches f inside Ω while the field matches g on the boundary. This is the classical statement, and it is the statement the finite-element method does *not* solve; it solves a reformulation that trades one derivative of the unknown for one derivative of a test function, and that trade is the entire reason the method works on the rough functions a computer can represent.

Definition 1.1 (Strong form of the Poisson problem). Given $f : \Omega \rightarrow \mathbb{R}$ and $g : \partial\Omega \rightarrow \mathbb{R}$, the strong form seeks $u : \bar{\Omega} \rightarrow \mathbb{R}$ with

$$-\nabla \cdot (\nabla u) = f \quad \text{in } \Omega, \quad u = g \quad \text{on } \partial\Omega. \quad (1.1)$$

A classical solution must be twice continuously differentiable, so that the Laplacian $\nabla \cdot (\nabla u)$ is defined pointwise.

The regularity demand of the strong form is the problem. A piecewise-linear function on a triangular mesh — the cheapest thing a finite-element code can build — is not twice differentiable; its second derivative is a sum of Dirac masses on the element edges. If we insist on the strong form, the discrete solution is illegal before we begin. The weak form removes the demand by integrating against a test function and moving one derivative across the integral.

To state it we need the spaces the derivatives are measured in. The natural space is not the space of smooth functions but the Sobolev space of functions whose first derivatives are square-integrable [3].

Definition 1.2 (Sobolev spaces H^1 and H_0^1). The space $H^1(\Omega)$ is the set of functions $v \in L^2(\Omega)$ whose weak gradient ∇v also lies in $L^2(\Omega)$, normed by

$$\|v\|_{H^1(\Omega)}^2 = \int_{\Omega} (v^2 + |\nabla v|^2) dx. \quad (1.2)$$

The subspace $H_0^1(\Omega)$ is the closure in this norm of the smooth functions with compact support in Ω ; its members vanish on $\partial\Omega$ in the trace sense. The trial space for the Poisson problem is H^1 , and the test space is H_0^1 .

The reason H^1 is exactly right, and not H^2 or C^2 , is that it is the largest space on which the weak form's bilinear form is finite and coercive. That is a claim we will prove, not assert. First the weak form itself.

Multiply the strong equation by a test function $v \in H_0^1(\Omega)$ and integrate over Ω . The left side carries a second derivative, which the divergence theorem moves onto v :

$$-\int_{\Omega} \nabla \cdot (\nabla u) v dx = \int_{\Omega} \nabla u \cdot \nabla v dx - \int_{\partial\Omega} (\nabla u \cdot n) v ds. \quad (1.3)$$

Because v vanishes on $\partial\Omega$, the boundary integral is zero, and one derivative has crossed from u to v . What remains is the weak form.

Theorem 1.3 (Weak form of the Poisson problem). A sufficiently smooth solution of the strong form satisfies, and is characterized by,

$$a(u, v) := \int_{\Omega} \nabla u \cdot \nabla v dx = \int_{\Omega} f v dx =: L(v) \quad \text{for all } v \in H_0^1(\Omega), \quad (1.4)$$

where a is a bilinear form on $H^1 \times H_0^1$ and L is a linear form on H_0^1 . This bilinear-and-linear pair is the object a finite-element library assembles [4].

Proof. Let u solve the strong form and let $v \in H_0^1(\Omega)$. Multiplying $-\nabla \cdot (\nabla u) = f$ by v and integrating gives $-\int_{\Omega} \nabla \cdot (\nabla u) v dx = \int_{\Omega} f v dx$. Applying the integration-by-parts identity above and using $v|_{\partial\Omega} = 0$ discards the boundary term, leaving $a(u, v) = L(v)$. Conversely, suppose $u \in H^2(\Omega)$ satisfies $a(u, v) = L(v)$ for every $v \in H_0^1$. Integrating the gradient term back by parts returns $\int_{\Omega} (-\nabla \cdot (\nabla u) - f)v dx = 0$ for all $v \in H_0^1$. Since the smooth compactly supported functions are dense in $L^2(\Omega)$ and the bracket lies in L^2 , the fundamental lemma of the calculus of variations forces $-\nabla \cdot (\nabla u) = f$ almost everywhere, recovering the strong equation. The two formulations therefore coincide on H^2 . \square

Bilinearity of a and linearity of L are immediate from the linearity of the integral and the gradient. What is not immediate — and what guarantees

Chapter 3

Boundary Conditions, Done Right: Dirichlet, Neumann, and Robin, Each Verified

Chapter 2 verified a solver on a problem whose boundary was pinned everywhere to a known value. Most problems are not so obliging: a heat sink holds a temperature on one face, a flux crosses another, and a convective law couples value and flux on a third. Each of those is a boundary condition of a different kind, each enters the finite-element method by a different mechanism, and each is a separate opportunity to be wrong in a way no picture reveals. This chapter treats the three canonical conditions — Dirichlet, Neumann, Robin — derives where each one enters the weak form, implements each in DOLFINx, and, following the discipline of the previous chapter, verifies each by a manufactured solution before trusting it.

The organizing idea is a single act of integration by parts. Moving one derivative off the trial function and onto the test function produces a boundary integral, and every boundary condition is a statement about that integral. A Dirichlet condition constrains the value and is imposed on the function space, so the boundary integral never appears; a Neumann condition prescribes the flux and *is* the boundary integral, added to the linear form; a Robin condition mixes the two and contributes to both forms. That

dichotomy — essential conditions on the space, natural conditions in the form — is the whole chapter in one sentence, and the sections that follow make it precise, prove that each implementation is well posed, and measure the order each one delivers. The deliverable is every condition implemented and individually verified, a mixed Dirichlet–Neumann–Robin problem confirmed to the a-priori order, and a missing boundary term caught as a collapsed order and a corrupted error field.

3.1 Essential and Natural: The Two Kinds of Boundary Condition

A boundary condition is a constraint the solution must satisfy on $\partial\Omega$, but the finite-element method sees two structurally different kinds, and the difference is not a matter of taste — it is dictated by where the condition can be enforced. One kind constrains the solution’s value and can only be imposed on the space of admissible functions; the other constrains the solution’s flux and can only enter through the weak form. Which is which falls out of integration by parts.

Definition 3.1 (Essential boundary condition). An essential boundary condition constrains the value of the solution on the boundary, $u = g$ on Γ_D . It is imposed directly on the trial and test spaces — the trial functions are required to take the value g and the test functions to vanish — and never appears as a term in the weak form.

Definition 3.2 (Natural boundary condition). A natural boundary condition constrains the flux of the solution — its outward normal derivative — on the boundary. It enters through a boundary integral in the weak form rather than through the space, so the trial and test spaces are unchanged; UFL represents it as an integral over the boundary measure [8].

Why the two kinds are handled so differently is not a convention but a consequence of the weak form’s structure. The boundary integral that carries the natural data appears only after integration by parts, and it is worth deriving explicitly because every later section substitutes into it.

Theorem 3.3 (The boundary term). Integrating the Poisson operator by parts against a test function v produces

$$\int_{\Omega} (-\nabla \cdot \nabla u) v \, dx = \int_{\Omega} \nabla u \cdot \nabla v \, dx - \int_{\partial\Omega} (\nabla u \cdot n) v \, ds, \quad (3.1)$$

so the weak form carries a boundary integral of the normal derivative $\partial_n u := \nabla u \cdot n$ times the test function. This boundary term is the exact location where natural conditions are substituted.

Proof. Apply the divergence theorem to the vector field $v \nabla u$: $\int_{\Omega} \nabla \cdot (v \nabla u) dx = \int_{\partial\Omega} v \nabla u \cdot n ds$. Expanding the divergence, $\nabla \cdot (v \nabla u) = \nabla v \cdot \nabla u + v \nabla \cdot \nabla u$. Substituting and rearranging, $\int_{\Omega} v \nabla \cdot \nabla u dx = \int_{\partial\Omega} v \partial_n u ds - \int_{\Omega} \nabla u \cdot \nabla v dx$. Negating both sides gives the claim. The manipulation requires $u \in H^2$ for a classical reading, but the identity extends to H^1 with the normal derivative interpreted in the weak trace sense [7]. □

The boundary term is the fork in the road. If we constrain v to vanish on part of the boundary, the term drops there and the value of u must be supplied some other way — on the space. If we leave v free and substitute a known value for $\partial_n u$, the term survives and carries the flux data into the form. That is the dichotomy, stated as a proposition because it governs every implementation choice in the chapter.

Proposition 3.4 (The essential/natural dichotomy). Every boundary condition for the Poisson problem is handled in exactly one of two ways: an essential condition shrinks the trial space — the test space loses the corresponding boundary degrees of freedom, so the boundary term vanishes there — while a natural condition modifies the bilinear or linear form by substituting known flux data into the boundary term. The two mechanisms are exclusive on any given piece of boundary.

Proof. On a piece of boundary where $u = g$ is imposed essentially, the test space is H_0^1 there, so $v = 0$ and the boundary integral $\int (\partial_n u) v ds$ vanishes identically; the value g is carried by the trial space, not the form. On a piece where the flux $\partial_n u$ is prescribed, v is unconstrained, the boundary integral does not vanish, and substituting the known $\partial_n u$ moves the data into the form. A single boundary piece cannot be both: prescribing $v = 0$ there kills the integral that a natural condition needs, and leaving v free there leaves the value g unenforced. Hence the two mechanisms partition the boundary [13]. □

Both mechanisms require the boundary to be regular enough that a value and a normal flux are defined at all — a mild condition met by every polygonal mesh.

A finite-element simulation you cannot verify is an opinion, not a result. This book solves real PDEs with DOLFINx and the FEniCSx stack and verifies every method it presents — in the open, by a test the reader can run — because verification comes before validation, and a slope you can measure beats a picture you can only admire.

The code is the argument, and all of it is here. Every figure comes from one self-contained Python program — the weak form in UFL, the mesh in gmsh, the system through PETSc, run in parallel over MPI — each printed in the book, runnable against a pinned stack.

Newton converges quadratically on an automatic Jacobian; a Stokes flow stays stable on an inf-sup pair; a steel wrench is solved from CAD as a coupled thermo-mechanical system.

Nothing is a black box, nothing left as an exercise.

Where a problem has a closed form, the answer is matched against it; where it does not, manufactured solutions measure the error against the order the theory promises. The book is candid about where finite elements quietly fail — volumetric locking, checkerboard pressure, geometry error hiding in a functional — and names the test that catches each. It builds to one verified analysis: a steel wrench, heated until its steel softens and loaded until it yields, reported down to the safety factor below one.

