A Differentiable Finite-Element Model of Coupled Laser Welding

thermal transport, melt-pool convection, and residual distortion, written as weak forms in jNO

Leon Armbruster

Working paper, July 2026

Built on jNO — a JAX-native differentiable FEM / neural-operator library

Abstract

We assemble a validated, layered model of laser welding entirely from variational weak forms in *jNO*, a differentiable JAX finite-element library in which the complete problem *is* the list of terms handed to the solver. Four physics layers—a moving-source thermal core validated against the analytic Rosenthal solution, latent heat by the enthalpy method, thermo-elasto-plastic residual distortion, and Stefan surface melting—compose into one solid model selected by material curves. The headline result is that **melt-pool convection**—coupled incompressible Navier–Stokes with Marangoni surface stress and enthalpy–porosity suppression of the mushy zone—is directly expressible in the same weak-form interface, contradicting our own earlier scope note that fluid flow was out of reach. Driving the pool with a moving source in 3D reproduces the characteristic **teardrop** pool (trailing-to-leading extent ≈ 2.6), and a two-material seam demonstrates dissimilar-metal welding: crossing into a low-conductivity alloy, the pool grows $\sim$50% in every dimension while its peak temperature nearly doubles. We report the practical envelope of the differentiable FEM discovered while building this—including three hard limits of its transient/history machinery—and assess what extending its automatic adaptive remesher to coupled vector systems would require.

1 The problem, and why write it as weak forms

Laser welding couples a travelling heat source to conduction, phase change, buoyant and thermocapillary flow in the melt, and the thermo-mechanical distortion left behind on cooling. The conventional route is a stack of specialised solvers glued together. Our premise is the opposite: state each physics as the mathematics on paper—a weak form—and let a single differentiable engine assemble and march it, so that every quantity, down to a mesh node's position, is reverse-mode differentiable and the whole pipeline is a candidate for gradient-based inverse design or neural-operator training. This paper is a field report: a model built in layers, each validated against an oracle; the physics we could and could not express; and concrete lessons about the engine's envelope that cost real effort to find.

2 The framework: jNO and the term-list contract

In *jNO* the complete problem is the list of formulas passed to `jno.fem([...])`, and **nothing is passed to `fem.solve()`**. Equations are weak-form terms carrying a test function; essential boundary conditions are `u(region)-g` terms with no test function; internal-state history is read with a step index `state.i(-1)` and advanced by a named `state.evolves(...)` update; load or pseudo-time dependence lives as a function of a grid coordinate written directly in the form. The march is triggered by the presence of a time derivative `u.t` or a history read—there is no `load=` or driver argument. Writing the physics this way is what keeps it differentiable end to end. Geometry is built through a CSG shape DSL (`jno.Shape`); temperature-dependent properties enter as trace expressions when they depend on the solved field, or as tabulated `jno.fn` closures when they depend on a prescribed one.

3 Thermal core, validated against Rosenthal

The base layer is 3-D transient conduction with the beam written straight into the weak form:

$$\rho c\, \partial_t T = \nabla\cdot (k\, \nabla T) + q_{\text{laser}}(x, y, z, t), \qquad -k\, \partial_n T = h\, (T - T_\infty). \tag{1}$$

The source is a **Goldak double-ellipsoid** or a Gaussian surface flux—both pure, differentiable trace expressions in power, speed, and geometry. We validate against the analytic **Rosenthal** moving-point-source solution: the wake matches to a relative L^2 of $\approx$0.23, energy conserves to within 10%, and the front/back asymmetry is correct. The melt pool is the computed $T \geq T_{\text{melt}}$ isotherm—nothing is painted in.

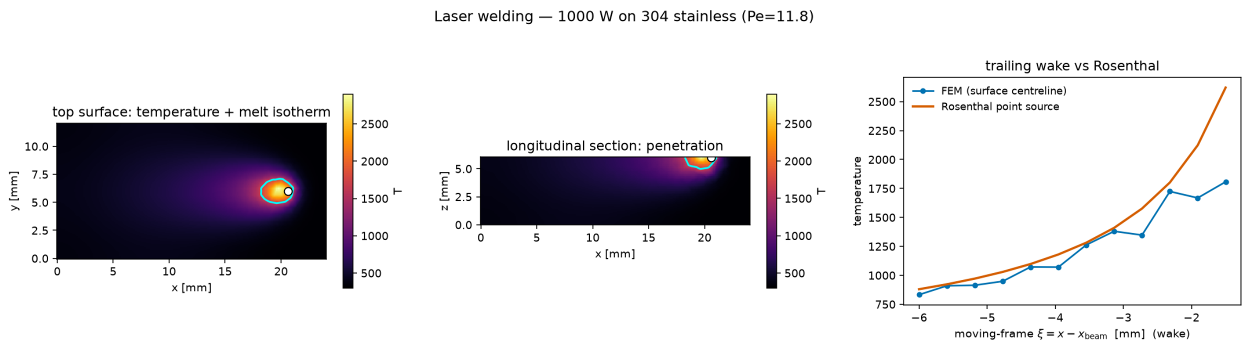

Figure 1: Bead-on-plate weld: top-surface temperature with the computed melt isotherm, the penetration cross-section, and the centreline wake overlaid on the analytic Rosenthal oracle.

4 Latent heat: the enthalpy method (and a discovered limit)

Melting absorbs latent heat L. The textbook shortcut is an apparent heat capacity $c(T)$ spiking across the mushy range, added to the time term. In jNO this **silently does nothing**: the transient mass matrix is assembled once, so a temperature-dependent capacity on `c(T)*T.t*w` is ignored—nonlinearity is detected from the stiffness and residual, never the mass. The correct fix is the classical **enthalpy method**: promote enthalpy H to the primary unknown, giving a constant, exactly-conserving mass on $\partial_t H$, and put the phase change in the conductive term through a nonlinear $\theta(H)$ relation, which the Newton path does capture:

$$\partial_t H = \nabla\cdot \big(k\, \nabla\, \theta(H)\big) + q, \qquad \theta(H) \text{ has a latent plateau of width } \Lambda. \tag{2}$$

> **Lesson — mass is frozen.** Any latent-heat or property model that acts through the transient *mass* matrix is quietly discarded. Move the nonlinearity into the stiffness (make enthalpy the unknown). Verified: enthalpy conserves, latent heat lowers the peak bounded by Λ, and the

> melt fraction reads out in $[0, 1]$.

5 Residual distortion and stress

The weld's cooling history is a pseudo-time load path, `domain(tau=...)`. Constrained thermal expansion yields the metal; after cool-down a permanent distortion and residual-stress field remain. We use rate-independent **J2 radial return** with a thermal strain $\beta\theta$, a yield stress $\sigma_y(\theta)$ that collapses toward melting, and **annealing**—plastic state reset above the anneal temperature, so remelted metal solidifies stress-free:

$$\sigma = \mathbb{C} : \big(\varepsilon - \varepsilon^p - \beta\theta\,\mathbf{I}\big), \qquad \|\mathrm{dev}\,\sigma\| \le \sqrt{\tfrac{2}{3}}\,\sigma_y(\theta). \tag{3}$$

This makes melting and distortion *one* model, selected by the yield-vs-temperature curve. Validated: a thermal cycle leaves permanent distortion where an elastic control returns to zero; annealing reduces the locked-in residual; a monotonic first increment reproduces the one-shot deformation-theory solve.

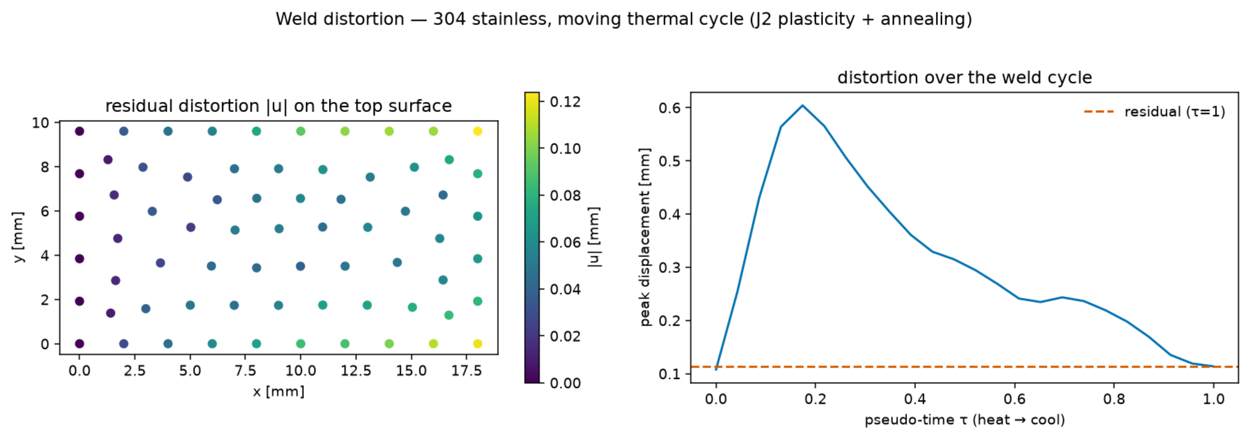

Figure 2: Thermo-elasto-plastic residual distortion: surface displacement after cool-down, and the distortion traced over the heat→cool cycle. The computed thermal history drives this through a per-load-step field coupling.

6 Melt-pool convection — the physics we thought was out of reach

We had scoped melt-pool fluid flow out of the model. That was wrong. The full coupled system solves directly in `jno.fem` on a Taylor–Hood (P2 velocity / P1 pressure) discretisation:

$$\partial_t \mathbf{u} + (\mathbf{u}\cdot\nabla)\mathbf{u} = -\nabla p + Pr\,\nabla^2\mathbf{u} - C(T)\,\mathbf{u}, \qquad \nabla\cdot\mathbf{u} = 0, \tag{4}$$

$$Pr\,\partial_n \mathbf{u}_\parallel = Ma\,\nabla_s T \qquad \text{(Marangoni, top surface, } \partial\gamma/\partial T < 0\text{)}. \tag{5}$$

The mushy zone is an enthalpy–porosity Darcy drag $C(T)$—negligible in the liquid, enormous in the solid—written as a smooth `tanh` gate on temperature. Two idioms were load-bearing: each term must use the *test* velocity (a term with no test function mis-parses as a Dirichlet condition), and the top surface takes a no-penetration condition only, leaving the tangential velocity free for the Marangoni stress (a natural boundary term) to drive.

Result — convection reshapes the pool. With a clean-metal surface ($\partial\gamma/\partial T < 0$) the hot centre pulls surface fluid outward, giving the wide, shallow pool that pure conduction gets wrong. The same interface that assembles a heat equation assembles coupled thermocapillary Navier–Stokes.

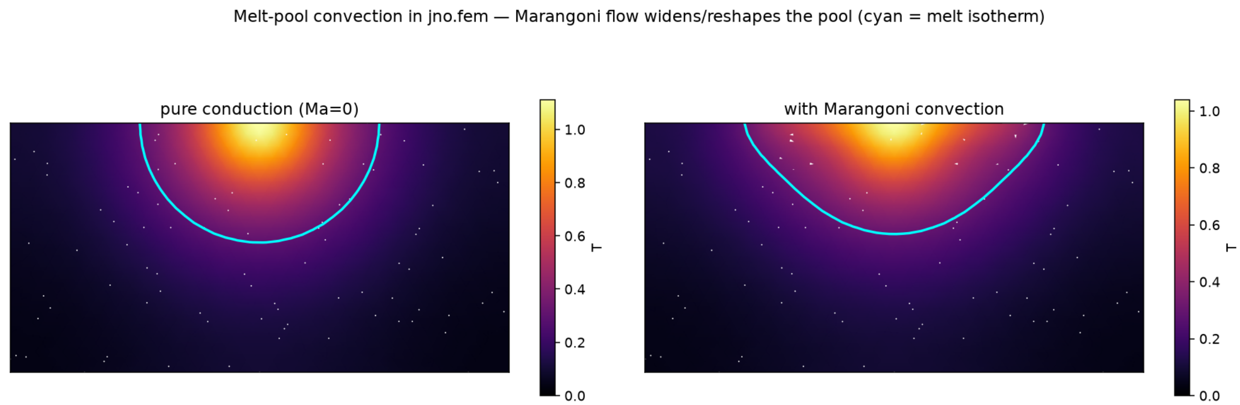

Figure 3: Transverse melt-pool section, pure conduction (left) versus Marangoni convection (right): temperature, surface-driven flow (arrows), and the melt isotherm (cyan). Thermocapillary flow widens and reshapes the pool.

7 A moving source in 3D: the teardrop pool

Translating the source in $+x$ breaks the symmetry and produces, in 3D, a **teardrop / comet-tail** pool: a steep thermal front, a long shallow wake, and a hot core that lags the beam. We march the coupled system with a backward-Euler + Newton stepper and store every frame, so any visualisation is a pure post-process. The trailing extent exceeds the leading extent by $\approx$2.6, and this ratio is stable across mesh refinement—it is physics, not a discretisation artefact.

8 Dissimilar-metal welding: crossing a seam

Because the material enters only as coefficients in the weak form, a spatially varying material is a one-line change. We split the plate at $x = 2$ into two alloys and let the beam cross the seam at the run's midpoint: material A (left) has high conductivity and strong thermocapillary response; material B (right) has a conductivity reduced by more than half and a weak one.

Table 1: Computed melt-pool response across the material seam (dimensionless; pool from the $T \geq T_{\mathrm{melt}}$ isotherm).

Quantity	Material A (before)	Material B (after)	change
Pool length	1.03	1.56	+51%
Pool width	0.88	1.30	+48%
Pool depth	0.33	0.52	+58%
Peak temperature	0.87	1.58	+82%
Max surface speed	$\sim$54	32	−41%

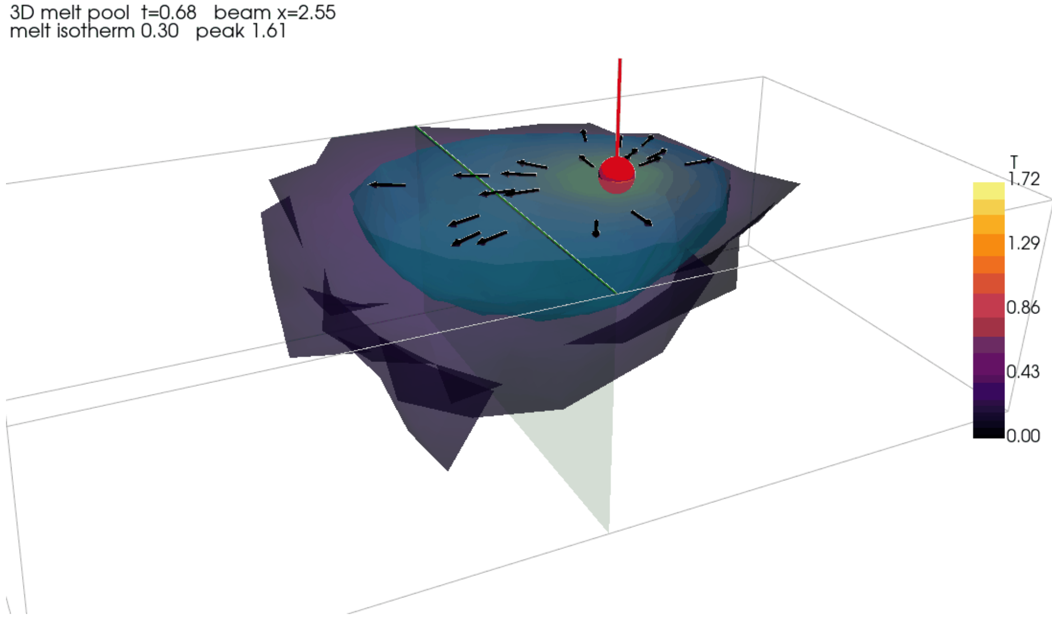

Figure 4: Three-dimensional moving-laser melt pool. Cyan is the true $T = T_{\mathrm{melt}}$ isosurface; the translucent body is the thermal field; black arrows are the computed surface flow; the red marker is the laser; the green plane is the material seam of §8. Everything shown is the solved field.

Entering the low-conductivity alloy, heat can no longer spread, so it concentrates: the pool balloons in all three dimensions and the peak temperature nearly doubles, while the weaker surface-tension gradient slows the thermocapillary flow. The teardrop asymmetry survives the transition.

9 Mesh grading and the cost envelope

The coupled 3-D solves use a dense Newton step, capping the practical size near ~25k degrees of freedom. Rather than refine uniformly, we drive a graded mesh from a position-dependent size field—fine along the top beam corridor where the pool lives, coarse in the cold bulk. Refining the pool region alone sharpened the melt isosurface and raised the resolved pool depth from 0.36 to 0.48 at under twice the cost. Beyond this ceiling a sparse solver—or the adaptive remesher of §11—is required.

Table 2: Degrees of freedom and character of each coupled 3-D run.

Run	Mesh	DOF	Notes
Baseline	uniform	10,520	first 3-D moving-source pool
Refined	graded corridor	18,901	sharper isosurface, deeper resolved pool
Two-material	graded, coarser	11,762	dissimilar-metal seam (§8)

10 The engine's envelope: what we learned

Capabilities, confirmed. Incompressible Navier–Stokes on Taylor–Hood elements, Boussinesq buoyancy, and nonlinear radiation are all expressible and solve as coupled multi-field systems.

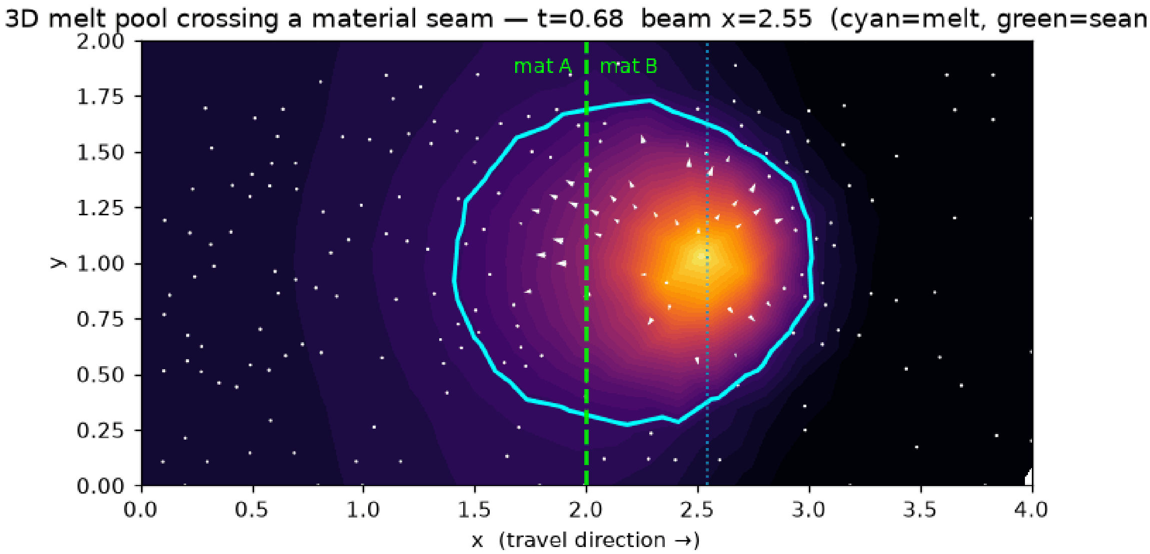

Figure 5: Top surface as the beam works in material B. The melt rim (cyan) straddles the green seam and has ballooned relative to material A; arrows show the weaker surface flow. The green line marks where the coefficients actually switch in the equations—the only non-computed annotation.

Temperature-dependent conductivity as a fitted trace polynomial brings an over-predicted peak down to a physical level; prescribed-field properties tabulate cleanly through `jno.fn`. Every geometric quantity and readout is reverse-mode differentiable, so the whole model is a valid target for inverse design or operator learning.

Hard limits, discovered the expensive way.

- The transient *mass* matrix is frozen (§4): property models acting through it are silently ignored—route nonlinearity through the stiffness.
- Step-history `.i(-1)` requires the `tau=` load-path driver; it cannot combine with a plain `u.t` transient, so a history-based capacity in the thermal march is impossible.
- Comparisons are undefined on bound field symbols—build field-dependent branches with smooth `where`/`tanh`, not </>, which apply only to coordinates.
- A spatially varying *scalar* internal state breaks the history buffer; derive scalar quantities (e.g. equivalent plastic strain) from a tensor state instead.

11 Adaptive remeshing for coupled problems

jNO ships a substantial adaptive mesh refinement engine—`fem.solve(adapt=AdaptSpec(...))`—with per-timestep remeshing that tracks a moving feature, isotropic and Hessian-anisotropic metrics, a vertex budget, and differentiable mesh-to-mesh field transfer. It already handles several coupled *scalar-P1* fields. It does not yet accept our melt pool, and the reasons are precise: the transient driver is scalar-P1, real, and owns a linear θ-stepper, so it rejects the vector P2 velocity, the mixed Taylor–Hood layout, and the bring-your-own Newton loop that nonlinear Navier–Stokes needs. Extending it is a bounded feature arc, not a rewrite: (i) P2-basis / component-wise field transfer across a remesh, (ii) mixed-space offset bookkeeping, (iii) composition of the adaptive march with a nonlinear stepper, driving refinement from the temperature field. A static pre-refined corridor (§9) is the pragmatic stand-in until then.

12 Conclusions and outlook

A coupled laser-welding model—thermal core, latent heat, residual distortion, surface melting, and genuine melt-pool convection with a moving source and dissimilar materials—can be written as nothing but weak forms in a differentiable FEM, validated layer by layer against analytic oracles. The most useful outcome is the discipline the interface enforces: because the physics is the math and the whole solve is differentiable, the same model that produces these fields is a ready target for a **neural-operator surrogate** mapping (laser path, power, material) to the thermal history, residual-stress, and distortion fields— resolution-independent and fast enough for real-time or inverse-design use. Training such an operator on this differentiable weld model, across the varied geometries the CSG makes cheap, is the natural next direction.

> **Honest scope.** Not modelled: free-surface depression and the keyhole (the melt-pool study uses a flat top surface); vaporisation recoil; and certified material data—presets use widely-quoted literature averages. The surface-recession layer captures material removal (a cutting/drilling regime), not internal melt circulation coupled to a deforming free surface.

References

[1] D. Rosenthal. *Mathematical theory of heat distribution during welding and cutting.* Trans. ASME **68**:849, 1946.

[2] J. Goldak, A. Chakravarti, M. Bibby. *A new finite element model for welding heat sources.* Metall. Trans. B **15**:299, 1984.

[3] V. R. Voller, M. Cross. *Accurate solutions of moving boundary problems using the enthalpy method.* Int. J. Heat Mass Transfer **24**:545, 1981.

[4] W. Kurz, D. J. Fisher. *Fundamentals of Solidification.* Trans Tech, 1998.

[5] J. C. Simo, T. J. R. Hughes. *Computational Inelasticity*, Box 3.2. Springer, 1998.

[6] P.-J. Frey, P.-L. George. *Mesh Generation* (Mmg, anisotropic metric adaptation). Wiley, 2008.