WHEN FULL-SPACE AND REDUCED-SPACE METHODS DIVERGE IN EQUALITY-CONSTRAINED OPTIMIZATION: A CONDITIONING PERSPECTIVE

Anonymous authors
Paper under double-blind review

ABSTRACT

Equality-constrained optimization is central to modern physical AI, where desired feasible solutions are typically computed via Newton-type methods due to their rapid convergence. However, when selecting between the two primary formulations, reduced-space vs. full-space, practitioners often rely on domain-specific heuristics rather than a principled framework. In this work, we show that the practical boundary between these two approaches is fundamentally governed by the spectral profile of the state constraint Jacobian. While both approaches succeed in well-conditioned settings, their numerical behaviors diverge under ill-conditioning. Our theoretical analysis shows that reduced-space methods excel when ill-conditioning is driven by excessively large singular values, whereas full-space methods prove superior when driven by vanishing singular values. Empirical evaluations across four diverse benchmarks validate our theoretical findings.

1 INTRODUCTION

Constrained optimization constitutes a foundational computational primitive in modern machine learning, physical AI, and autonomous control (Andersson et al., 2015; Amos et al., 2018; Kotary et al., 2021). Across applications—ranging from enforcing exact physical conservation laws in physics-informed neural networks (PINNs) (Song et al., 2024; Luo et al., 2025) to computing implicit gradients in deep equilibrium models (Bai et al., 2019; 2021) and generating dynamic trajectories in real-time model predictive control (Ghazaei Ardakani et al., 2019; Luis et al., 2020), high-dimensional programs demand precise and reliable solutions. Equality constraints play a pivotal role in these problems by embedding fundamental physical laws that must be strictly satisfied.

Generally, an equality-constrained optimization problem takes the following form:

$$\underset{u\in\mathbb{R}^n}{\text{minimize}}\ f(x,u), \quad \text{subject to } g(x,u)=0, \tag{1}$$

where $u \in \mathbb{R}^n$ denotes design variables which are directly optimized, and $x \in \mathbb{R}^m$ represents state variables implicitly determined by u. State-of-the-art solution algorithms broadly divide into two distinct paradigms: full-space and reduced-space formulations. Reduced-space methods leverage the implicit function theorem to eliminate the state variables and obtain the mapping $x = x(u)$ via the equality constraints, thereby restricting the optimization search trajectory strictly to the lower-dimensional constraint manifold. Conversely, full-space methods treat state variables, control variables, and Lagrange multipliers symmetrically by formulating the problem within a unified Lagrangian function

$$\mathcal{L}(z,\lambda) = f(z) + \lambda^T g(z) \tag{2}$$

where $z = [x, u]$ denotes the primal variables and λ represents the dual variables associated with the equality constraints, updating them simultaneously by solving large Karush–Kuhn–Tucker (KKT) systems. To focus on how complex constraints interact with numerical optimizers, we assume all constraints are equality constraints.

Currently, the choice between deploying a reduced-space or full-space framework is largely driven by empirical heuristics and implementation convenience (Biegler, 2010; Hicken & Alonso, 2013). For instance, reduced-space approaches are preferred when the number of design variables is significantly

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

smaller than the number of state variables, whereas full-space approaches are favored when the expanded KKT matrix structures present exploitable sparsity (Nocedal & Wright, 2006; Betts, 2010). While each scheme excels in specific domains, such as reduced-space methods in topology and shape optimization (Sigmund & Maute, 2013; Deaton & Grandhi, 2014) and full-space methods in dynamic and trajectory optimization (Hargraves & Paris, 1987; Biegler, 2007), there remains a lack of generalized understanding regarding the fundamental boundary separating these two paradigms, as existing comparative analyses remain predominantly domain-specific.

In this work, we analyze this practical boundary through the lens of numerical conditioning. Specifically, we show that the choice between full-space and reduced-space paradigms is fundamentally governed by the condition number of the state constraint Jacobian $\partial g/\partial x$. We demonstrate that both paradigms perform effectively when $\partial g/\partial x$ is well-conditioned. Under ill-conditioning, however, their numerical behavior diverges based on the underlying spectral structure: reduced-space methods are favored when ill-conditioning stems from excessively large singular values, whereas full-space methods excel when it is driven by vanishing singular values. We support this claim with a rigorous theoretical analysis, and validate our findings across several practical optimization benchmarks. In this paper, we evaluate four real-world engineering problems spanning neural PDEs, trajectory optimization, structural design optimization, and thermal topology optimization. Each problem is constructed to avoid trivial failure modes, such as breakdowns in auxiliary state solvers or prohibitively large KKT systems. These benchmarks are deliberately designed to illustrate that conventional heuristics based on problem scale or matrix sparsity cannot explain the observed performance gap. Instead, they highlight the central role of the spectral conditioning of the state constraint Jacobian.

The main contributions of this work are summarized as follows:

- We establish a theoretical framework demonstrating that the spectral profile of the state constraint Jacobian, $\partial g/\partial x$, serves as the fundamental criterion dictating the efficacy and stability of full-space versus reduced-space paradigms.
- We rigorously validate our theory across diverse engineering and optimization benchmarks, confirming that the singular value profile of $\partial g/\partial x$ consistently predicts solver performance in practice.
- Based on our theoretical and empirical insights, we propose a simple rule that practitioners can easily deploy to select the optimal solver formulation for a given problem instance.

2 Related Work

Learning-based optimization schemes. A prominent line of research focuses on optimization-layer methods that embed numerical solvers directly as differentiable layers within end-to-end architectures, as exemplified by OptNet (Amos & Kolter, 2017). These approaches typically adopt a full-space formulation and rely heavily on quadratic programming (QP) structures with quadratic objectives and static, linear constraints (Agrawal et al., 2019). However, while quadratic/convex objectives are common, real-world equality constraints that encode complex physical or domain-specific rules are often highly nonlinear (Sigmund & Maute, 2013; Deaton & Grandhi, 2014; Gould et al., 2022).

Another line of research focuses on handling more complex equality constraints through reduced-space formulations, where constraints are solved separately (Donti et al., 2021; Fioretto et al., 2021). For instance, tailored reduced-space methods have been developed to solve PINN-constrained optimization directly in the neural parameter space whose high dimensionality makes full-space approaches computationally intractable (Hao et al., 2023; Wang et al., 2025). However, these approaches inherit the inherent limitations of reduced-space frameworks, where repeated nonlinear constraint solves and adjoint sensitivity computations can become severe performance bottlenecks. Consequently, a systematic understanding of when to deploy reduced-space versus full-space formulations in optimization layers remains lacking.

Theoretical aspects of numerical optimization. A growing body of work in the machine learning community focuses on analyzing and enhancing Newton-type numerical optimization methods. For instance, recent studies refine Hessian approximations in quasi-Newton methods (Liu & Luo, 2022; Fang et al., 2025; Romanov et al., 2025; Wang & Qiu, 2025) and improve line-search strategies (Jin

& Mokhtari, 2021; Jin et al., 2024; Hanzely et al., 2026), while various algorithmic variants have been introduced to accelerate convergence (Lin et al., 2021; Jiang & Mokhtari, 2023).

Companion theoretical works have similarly established non-asymptotic convergence bounds for Newton variants (Doikov & Nesterov, 2020; Hanzely et al., 2022). However, the vast majority of these studies focus strictly on unconstrained optimization with convex objectives, a setting that fails to capture the complex computational dynamics, non-convexity, and manifold constraints inherent to practical equality-constrained problems.

3 PRELIMINARIES

Assumptions. While constructing a theoretical model that perfectly captures real-world complexity is intractable, we make the following assumptions to represent typical engineering scenarios:

1. *The objective function $f(z)$ is smooth, convex in the primal variables $z = [x, u] \in \mathbb{R}^{m+n}$, where the dimension of design variables is smaller than that of state variables ($n < m$).*
2. *The equality constraint $g(z)$ is a smooth function of $z = [x, u] \in \mathbb{R}^{m+n}$, and the Jacobian $\partial g/\partial x$ remains invertible so that the implicit mapping $x = x(u)$ exists at points where $g(x, u) = 0$.*
3. *The equality constraint Jacobian with respect to design variables u, namely $\partial g/\partial u$, is well-conditioned. We also assume that the Hessian $\nabla^2_{zz} f(z)$ is well-conditioned.*

Note that the assumptions on the objective function f reflect a general case, as f is entirely user-defined. In contrast, the reduced objective $f(x(u), u)$, where $x(u)$ represents the implicit mapping defined by $g(x, u) = 0$, is not guaranteed to be convex with respect to $u \in \mathbb{R}^n$. Furthermore, we assume all functions are sufficiently smooth, ensuring that numerical solvers do not fail due to simple irregularity and allowing our analysis to focus strictly on optimization dynamics.

Notation. We denote $\partial g/\partial x$, $\partial g/\partial u$ and $\partial g/\partial z$ by J_x, J_u and J_z, respectively. For a matrix A, $\sigma_{\max}(A)$ and $\sigma_{\min}(A)$ represent its largest and smallest singular value, respectively, where $\|Av\|$ is the Euclidean norm of the vector Av. $\|A\| = \sigma_{\max}(A)$ represents its 2-norm (a.k.a, spectral norm). The condition number of A is denoted by $\kappa(A) = \sigma_{\max}(A)/\sigma_{\min}(A)$. We also abbreviate the term full-space and reduced-space by FS and RS, respectively.

4 CONDITIONING ANALYSIS FOR NEWTON METHODS

In this section, we systematically analyze how the singular values of the equality constraint Jacobian J_x influence the conditioning of the Hessian matrix within each scheme. Consider the standard Newton iteration:

$$u_{k+1} = u_k - \alpha_k H_k^{-1} \nabla f(u_k) \tag{3}$$

where $\alpha_k \in (0, 1]$ represents the line-search step size and H_k is the Hessian matrix. As a linear system, computing the search direction $d_k = -H_k^{-1} \nabla f(u_k)$ is prone to severe numerical instability and cancellation errors when H_k is ill-conditioned. Although line search ensures sufficient decrease by backtracking on α_k, a highly inaccurate search direction forces $\alpha_k \to 0$, leading to premature termination or numerical stagnation. Consequently, modern solvers incorporate adaptive regularization; for example, IPOPT Wächter & Biegler (2006) modifies the Hessian to $H_k + \mu_k I$ with a dynamic parameter $\mu_k > 0$. When ill-conditioning degrades the system, these globalization mechanics aggressively increase μ_k, pushing the solver into the heavily damped regime where $\mu_k \gg 1$. Applying a Neumann series expansion to $(H_k + \mu_k I)^{-1}$ yields

$$(H_k + \mu_k I)^{-1} = \frac{1}{\mu_k} I - \frac{1}{\mu_k^2} H_k + \mathcal{O}\left(\mu_k^{-3}\right). \tag{4}$$

Substituting this expansion into the search direction yields

$$d_k^u = -\frac{1}{\mu_k} \nabla f(u_k) + \frac{1}{\mu_k^2} H_k \nabla f(u_k) + \mathcal{O}\left(\frac{\|H_k\|^2 \|\nabla f(u_k)\|}{\mu_k^3}\right) \tag{5}$$

revealing that aggressive growth in μ_k directly suppresses the step magnitude in a manner analogous to line-search backtracking. As illustrated in Figure 1, both regularization and line search inevitably force small steps, severely hindering the overall convergence rate and computational efficiency of the Newton method.

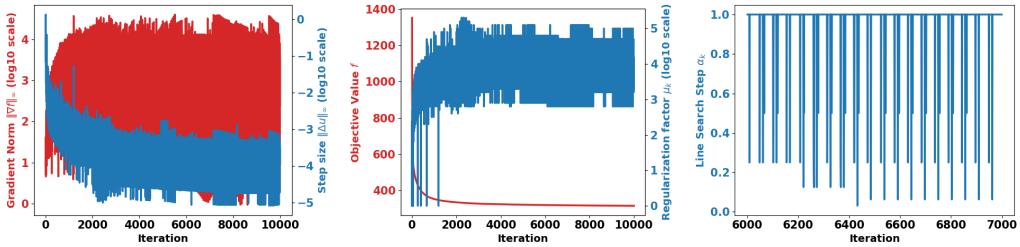


Figure 1: Illustration of the training dynamics of the Newton method when Hessian becomes ill-conditioned

Inverted Jacobian in RS methods. Thus, whether an optimization problem can be solved efficiently hinges on the conditioning of the Newton Hessian system. This section investigates the extent to which equality constraints affect this conditioning. Specifically, RS methods optimize the reduced objective $f(x(u), u)$ with respect to u via the total gradient

$$\nabla_u f(x(u), u) = \frac{\partial f}{\partial x}\frac{\mathrm{d}x}{\mathrm{d}u} + \frac{\partial f}{\partial u} = \frac{\partial f}{\partial x} J_x^{-1} J_u + \frac{\partial f}{\partial u}$$

which implies that the conditioning of the matrix $J_x^{-1}J_u$ is critical to the efficacy of RS methods. To examine how J_u influences its spectrum, let $v \in \mathbb{R}^n$ with $\|v\| = 1$, and consider the norm $\|J_x^{-1}J_u v\|$. If J_x^{-1} is ill-conditioned, then either $\sigma_{\min}(J_x^{-1}) \ll 1$ or $\sigma_{\max}(J_x^{-1}) \gg 1$. Consequently, the magnitude of $\|J_x^{-1}J_u v\|$ depends heavily on the extent to which $J_u v$ projects onto these ill-conditioned directions.

Assume further that $v \mapsto J_u v$ is agnostic to J_x^{-1}, so that $J_u v$ has balanced projections across all right singular vectors of J_x^{-1}. Consequently, $\|J_x^{-1}J_u v\|$ is bounded away from zero even when $\sigma_{\min}(J_x^{-1}) \ll 1$, as non-vanishing singular values prevent the norm from decaying. On the other hand, if $\sigma_{\max}(J_x^{-1}) \gg 1$, a non-trivial projection onto the corresponding singular vector can significantly amplify $\|J_x^{-1}J_u v\|$. We formalize this behavior in the following theorem:

Theorem 4.1. *Let $\{w_i\}_{i=1}^m$ denote the orthonormal right singular vectors of J_x^{-1}, and assume that there exists $c > 0$ such that $|w_i^T J_u v| \geq c\|J_u v\|$ for all $i = 1, ..., m$ and $v \in \mathbb{R}^n$. Then we have*

$$c\,\sigma_{\min}(J_u)\,\sigma_{\max}(J_x^{-1}) \leq \|J_x^{-1}J_u v\| \leq \sigma_{\max}(J_u)\,\sigma_{\max}(J_x^{-1})$$

where $\|v\| = 1$.

Thus, RS methods are generally well-conditioned as long as $\sigma_{\max}(J_x^{-1})$ remains bounded, which is equivalent to $\sigma_{\min}(J_x)$ staying bounded away from zero. Conversely, if $\sigma_{\min}(J_x) \to 0$, both the reduced gradient computations and the reduced Hessian matrix become severely ill-conditioned, degrading performance and forcing line-search or regularization mechanisms to stall.

Augmented Jacobian in FS methods. Full-space methods iteratively solve the following KKT system

$$\begin{bmatrix} B_k & \frac{\partial g}{\partial z}(z_k)^T \\ \frac{\partial g}{\partial z}(z_k) & 0 \end{bmatrix} \begin{bmatrix} d_k^z \\ d_k^\lambda \end{bmatrix} = \begin{bmatrix} \nabla_z \mathcal{L}(z_k, \lambda_k) \\ g(z_k) \end{bmatrix} \tag{6}$$

to update the primal-dual variables $[z_{k+1}, \lambda_{k+1}] = [z_k, \lambda_k] + [d_k^z, d_k^\lambda]$. Here, $B_k = \nabla_{zz}^2 \mathcal{L}(z_k, \lambda_k)$ denotes the Hessian of the Lagrangian $\mathcal{L}$ for exact Newton methods, or a quasi-Newton approximation (e.g., BFGS). In practice, positive-definite approximations are often used when exact Hessians are computationally prohibitive. A notable special case arises in equality-constrained QP, such as the equality-only setting of OptNet, where $f(z) = \frac{1}{2}z^T Q z + q^T z$ and $g(z) = Az - b$. In this case, the KKT system equation 6 simplifies to the linear system

$$\begin{bmatrix} Q & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} d_k^z \\ d_k^\lambda \end{bmatrix} = \begin{bmatrix} Qz_k + q + A^T z_k \\ Az_k - b \end{bmatrix}. \tag{7}$$

Assuming the saddle-point matrix is nonsingular, the exact optimal primal-dual solution $[z^*, \lambda^*]$ is obtained in a single step directly from this linear system, bypassing the need for iterative Newton updates. The following theorem demonstrates that the singular value spectrum of the equality constraint Jacobian directly governs the conditioning of the KKT matrix:

Theorem 4.2. *Let* $W = \begin{bmatrix} B & J_z^T \\ J_z & 0 \end{bmatrix}$ *be a KKT matrix. Assume that* J_z *is full row rank and* $\|v^T B v\| > 0$ *for all* $v \in \ker(J_z) \backslash \{\mathbf{0}\}$*, then we have*

$$\kappa_2(W) \leq \frac{\mu_{\max} + \sqrt{\mu_{\max}^2 + 4\sigma_{\max}(J_z)^2}}{2 \cdot \min(\frac{\sqrt{\mu_{\max}^2 + 4\sigma_{\min}(J_z)^2} - \mu_{\max}}{2}, \mu_{\min})} \tag{8}$$

where $\mu_{\max} = \sigma_{\max}(B)$ *and* $\mu_{\min} = \sigma_{\min}(B)$.

Because B decouples state and design variables, it is typically well-conditioned in practice, leaving the conditioning of W primarily governed by the singular values of J_z. Note that the upper bound in Equation equation 8 becomes large when $\sigma_{\max}(J_z) \gg 1$ or $\sigma_{\min}(J_z) \ll 1$. Because the Jacobian $J_z \in \mathbb{R}^{m \times (m+n)}$ has more columns than rows, its singular values are the non-negative square roots of the eigenvalues of $J_z J_z^T$, where $J_z (J_z)^T = [J_x, J_u][J_x, J_u]^T = J_x J_x^T + J_u J_u^T$. Applying Weyl's inequality yields

$$\sigma_{\max}(J_x) \leq \sigma_{\max}(J_z) \leq \sqrt{\sigma_{\max}^2(J_x) + \sigma_{\max}^2(J_u)}, \tag{9}$$

which implies that $\sigma_{\max}(J_z)$ remains bounded provided that $\sigma_{\max}(J_x)$ is not excessively large. Conversely, because $n < m$, the matrix $J_u J_u^T$ is rank-deficient and possesses zero eigenvalues; hence, a lower bound cannot be obtained directly via Weyl's inequality. However, for any unit vector $v \in \mathbb{R}^m$, we have $\|J_z^T v\|^2 = \|J_x^T v\|^2 + \|J_u^T v\|^2$. Consequently, $\|J_z^T v\|$ is small only if both $\|J_x^T v\|$ and $\|J_u^T v\|$ are small simultaneously. This requires the singular vectors of J_x^T corresponding to near-zero singular values to lie almost entirely within the $(m-n)$-dimensional subspace $\ker(J_u^T)$, a condition that is highly unlikely in practice. We make this intuition precise in the following theorem:

Theorem 4.3. *Let* $\{v_i\}_{i=1}^m$ *denote the orthonormal right singular vectors of* J_x^T *with the corresponding singular value* σ_i*. Given* $\delta \geq \sigma_{\min}(J_x^T)$*, let* V_δ *denote the subspace spanned by all right singular vectors of* J_x^T *whose singular value is no greater than* δ*. Assume that* $V_\delta \cap \ker(J_u^T) = \{\mathbf{0}\}$*, then we have*

$$\sigma_{\min}(J_z) \geq \min(\delta, \sigma_{\min}(J_u)) \cdot \sqrt{1 - \cos\theta},$$

where $\cos\theta = \max\{|v^T w| : v \in V_\delta, w \in \ker(J_u^T), \|v\| = \|w\| = 1\} < 1$.

Therefore, a well-conditioned J_u generally elevates the smallest singular value of J_z, but it cannot mitigate the overall ill-conditioning induced by an excessively large $\sigma_{\max}(J_x)$, as indicated in Equation 9.

Boundary between two schemes. The design constraint Jacobian J_u plays an analogous regularizing role in both frameworks, mitigating the numerical impact of adverse singular values in J_x. The fundamental distinction lies in the operator space where this regularization acts: J_u regularizes J_x directly in FS formulations, whereas it regularizes the inverse sensitivity J_x^{-1} in RS methods. Because $\kappa(J_x) = \kappa(J_x^{-1})$, both schemes yield well-conditioned system matrices when J_x is well-conditioned. However, when J_x becomes severely ill-conditioned, the specific failure mechanism diverges depending on the spectrum of J_x: an excessively large maximum singular value $\sigma_{\max}(J_x)$ degrades FS KKT conditioning, whereas a vanishing minimum singular value $\sigma_{\min}(J_x)$ causes the RS Hessian H_R to explode. Algorithm 1 is an example of a simple scheme-selection framework.

5 Experiments

We assess the performance of the RS and FS approaches on a suite of practical NLP benchmarks from different engineering domains. The underlying optimization problems are formulated via CasADi (Andersson et al., 2019) and solved using IPOPT (Wächter & Biegler, 2006) with default algorithmic hyperparameters and convergence tolerances unless otherwise specified. Full implementation specifications and reproducible problem formulations are provided in Appendix A.

Neural PDE optimization. We first consider a neural-PDE-constrained tracking problem:

$$\min_{\boldsymbol{v},\boldsymbol{u}} \quad \sum_{i=0}^{N} \Big[(v_i - \bar{v}_i)^2 + \gamma u_i^2 \Big], \quad \text{s.t.} \quad R(\boldsymbol{v}, \boldsymbol{u}) = 0$$

where the equality constraint $R(\boldsymbol{v}, \boldsymbol{u})$ is modeled by a neural network surrogate trained on the steady-state forced viscous Burgers' equation $vv_x - \nu v_{xx} = u(x)$ on a one-dimensional domain $x \in [0, 1]$ from PDEBench (Takamoto et al., 2022). Here, $v(x)$ represents the non-linear velocity state field, $\bar{v}(x)$ is the target profile, $\nu = 0.001$ is the kinematic viscosity coefficient, and $u(x)$ is a distributed spatial control source acting as the driving forcing function. We adopt the RS framework proposed by Hao et al. (2023). Direct optimization over the neural network weights is intentionally avoided, since it would result in an excessively high-dimensional state space ($\gtrsim 10^5$ parameters) that imposes severe memory and computational bottlenecks on FS KKT matrix factorizations.

Figure 2 compares the performance of the RS and FS formulations evaluated using IPOPT. The FS scheme robustly converges across all mesh resolutions, maintaining optimal objective values. In contrast, the final objective value achieved by the RS scheme systematically degrades as N increases, directly tracking the decay of $\sigma_{\min}(J_v)$. This empirical trend confirms our theoretical prediction: while the ill-conditioning of J_v is driven primarily by a vanishing minimum singular value ($\sigma_{\min}(J_v) \to 0$), this rank deficiency selectively destabilizes the RS formulation via the inverse sensitivity $J_v^{-1} J_u$, whereas the FS KKT framework remains resilient. All singular values are evaluated at the termination of the optimization process.

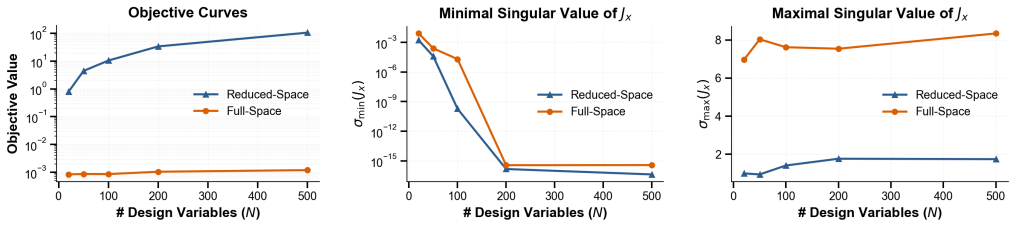


Figure 2: Training results for RS and FS methods in the neural PDE problem.

Trajectory optimization. Our second benchmark considers the classic underactuated double pendulum swing-up and balance problem. The control objective is to transition the system from its downward resting state to the unstable upright equilibrium at the origin while minimizing the cumulative quadratic state-control deviation. Formulated via direct transcription, the discrete actuation inputs u_k and state variables x_k (angular positions and velocities) are constrained by the discretized Euler–Lagrange dynamics $F(x_k, x_{k+1}, u_k) = 0$. The discrete-time optimal control problem is expressed as:

$$\min_{\boldsymbol{x},\boldsymbol{u}} \quad J(\boldsymbol{x}, \boldsymbol{u}) = \sum_{t=0}^{T-1} \Big[x_t^T Q x_t + u_t^T R u_t \Big] + \Psi(x_T), \quad \text{subject to} \quad x_{t+1} = f(x_t, u_t)$$

with initial state x_0, $\boldsymbol{u} = (u_0, ..., u_{T-1})$ is the control sequence where $f(x_t, u_t)$ is the dynamical model of the double pendulum system. The state space $x = [\theta_1, \theta_2, \dot{\theta}_1, \dot{\theta}_2]^T$ is four-dimensional and controls are one-dimensional. To eliminate potential errors from external residual solvers, we generate the state sequence explicitly forward through the discrete dynamics, bypassing state equation solves entirely.

Table 1 compares the performance of both optimization formulations across increasing time horizons T. As T extends, the condition number of the reduced Hessian $\kappa(H_R)$ degrades rapidly, directly tracking the decay of $\sigma_{\min}(J_x)$. In contrast, the full-space KKT matrix remains numerically robust despite the vanishing singular values of J_x. Note that experimental *success* is defined strictly by optimizer convergence rather than state trajectory quality; a run is deemed successful when IPOPT satisfies its default convergence criterion (an overall NLP error threshold of 10^{-8}). Runs exceeding 10000 iterations without meeting this tolerance are terminated and marked as failures.

Table 1: Performance of RS (top) and FS (bottom) methods in the trajectory optimization problem.

T	n_{var}	Success	Iters	Obj. value	$\sigma_{\min}(J_x)$	$\sigma_{\max}(J_x)$	$\kappa(H_{\text{R}})$
40	40	✓	4855	3.1×10^{2}	1.3×10^{-8}	28.2	3.7×10^{9}
50	50	×	10000	3.1×10^{2}	2.7×10^{-13}	29.3	1.2×10^{14}
100	100	×	10000	1.1×10^{3}	7.1×10^{-15}	32.1	3.3×10^{15}
200	200	×	10000	2.5×10^{3}	1.5×10^{-17}	33.1	7.5×10^{18}
500	500	×	10000	8.1×10^{3}	7.3×10^{-14}	32.7	1.4×10^{13}
T	n_{var}	Success	Iters	Obj. value	$\sigma_{\min}(J_x)$	$\sigma_{\max}(J_x)$	$\kappa(H_{\text{KKT}})$
50	254	✓	21	3.1×10^{2}	2.2×10^{-8}	10.3	2.7×10^{4}
100	504	✓	43	6.9×10^{2}	3.9×10^{-13}	35.2	1.0×10^{5}
200	1004	✓	89	6.9×10^{2}	2.5×10^{-15}	35.2	1.0×10^{5}
500	2504	✓	162	1.1×10^{3}	8.4×10^{-17}	35.2	1.0×10^{5}
1000	5004	✓	456	1.1×10^{3}	1.0×10^{-15}	35.2	1.0×10^{5}
2000	10004	✓	1036	1.1×10^{3}	2.4×10^{-17}	35.2	1.0×10^{5}

Structural design optimization. The third benchmark is from structural design optimization, a domain traditionally dominated by RS methods. We consider a one-dimensional, tip-loaded Euler–Bernoulli cantilever sizing problem: choose element heights $u = (u_1, \ldots, u_N)^\top$ to minimize structural compliance $c = f^\top x$ subject to discrete equilibrium and a fixed material volume. With N Hermite cubic elements the state is $x \in \mathbb{R}^{2N}$ (nodal deflections and rotations) and the residual is $R(x, u) = K(u)x - f$:

$$\min_{u,x} \; f^\top x \quad \text{s.t.} \quad R(x,u) = 0, \;\; bL_e \sum_{e=1}^{N} u_e = V_0, \tag{10}$$

where b is the out-of-plane beam width, L_e the element length, and V_0 the prescribed total volume. The strong form is the *fourth-order* bending equation $(EI\, w'')'' = q$, which after discretization yields $K_e(u) = \frac{Eb}{12L_e^3} u_e^3 \widehat{K}$ with height-independent template $\widehat{K}$. Thus $J_x = \partial R/\partial x = K(u)$ is a discrete fourth-order operator.

Table 2: Performance of RS (top) and FS (bottom) methods in the beam design optimization problem.

N	n_{var}	Success	Iters	Obj. value	$\sigma_{\min}(J_x)$	$\sigma_{\max}(J_x)$	$\kappa(H_{\text{R}})$
20	20	✓	29	3.42	1.5×10^{-1}	4.5×10^{4}	5.8×10^{0}
50	50	✓	25	3.4	6.3×10^{-2}	7.4×10^{5}	9.3×10^{0}
100	100	✓	30	3.39	3.2×10^{-2}	6×10^{6}	1.3×10^{1}
200	200	✓	33	3.39	1.6×10^{-2}	4.8×10^{7}	1.9×10^{1}
500	500	✓	49	3.39	6.4×10^{-3}	7.5×10^{8}	2.9×10^{1}
1000	1000	✓	75	3.39	3.2×10^{-3}	6×10^{9}	2.6×10^{1}
N	n_{var}	Success	Iters	Obj. value	$\sigma_{\min}(J_x)$	$\sigma_{\max}(J_x)$	$\kappa(H_{\text{KKT}})$
20	60	✓	27	3.42	2.2×10^{-1}	1.6×10^{6}	2.7×10^{8}
50	150	✓	31	3.4	9.4×10^{-2}	6.1×10^{7}	1.6×10^{10}
100	300	✓	36	3.39	4.8×10^{-2}	9.6×10^{8}	4×10^{11}
200	600	×	500	1.04	6.6×10^{-2}	1.3×10^{9}	1.8×10^{13}
500	1500	×	500	1.03	2.2×10^{-2}	5.7×10^{11}	1.6×10^{15}
1000	3000	×	500	1.04	1.1×10^{-2}	8.9×10^{12}	6.3×10^{15}

Table 2 reports scaling of RS and FS with mesh size N. RS converges through $N = 1000$ with compliance consistent with the closed-form optimum $c^\star \approx 3.39$. FS succeeds only through $N = 100$ and fails for $N \geq 200$. In this problem, the unphysically low objective values reported by the FS method at high resolutions come directly from a failure to satisfy static equilibrium constraints, resulting in infeasible solutions. As N increases, $\sigma_{\max}(J_x)$ grows rapidly, consistent with $\kappa(H_{\text{KKT}}) \sim \mathcal{O}(N^{4.54})$ when FS exposes J_x in the outer KKT system (Theorem 4.2). RS optimizes only in u, factors $K(u)$ inside state and adjoint solves, and tracks $\kappa(H_{\text{R}}) \sim \mathcal{O}(N^{0.41})$, so the reduced Hessian remains

Table 3: Performance of RS (top) and FS (bottom) methods in the thermal topology optimization problem.

N	n_{var}	Success	Iters	Obj. value	$\sigma_{\min}(J_x)$	$\sigma_{\max}(J_x)$	$\kappa(H_{\text{R}})$
50	50	✓	23	1.27	1.1×10^{-2}	9.4×10^{1}	4.9×10^{2}
100	100	✓	28	1.3	6.5×10^{-3}	2×10^{2}	8.6×10^{2}
200	200	✓	88	1.35	3.6×10^{-3}	4.3×10^{2}	4.1×10^{3}
400	400	✓	35	1.51	2.1×10^{-3}	9×10^{2}	3.2×10^{5}
500	500	✓	37	1.6	1.7×10^{-3}	1.1×10^{3}	2.9×10^{6}
N	n_{var}	Success	Iters	Obj. value	$\sigma_{\min}(J_x)$	$\sigma_{\max}(J_x)$	$\kappa(H_{\text{KKT}})$
50	99	✓	346	1.29	1.4×10^{-2}	1×10^{2}	3.1×10^{5}
100	199	✓	81	1.3	6.5×10^{-3}	2×10^{2}	5.1×10^{6}
200	399	✓	89	1.35	3.7×10^{-3}	4.3×10^{2}	4.4×10^{7}
400	799	✓	92	1.52	1.9×10^{-3}	9.3×10^{2}	4.9×10^{8}
500	999	✓	179	1.6	1.6×10^{-3}	1.1×10^{3}	5.1×10^{8}

comparatively well conditioned even when $\sigma_{\max}(J_x) \gg 1$. Figure 3 shows that ill-conditioning of H_{KKT} limits FS on fine meshes long before the physics changes, whereas RS delays that bottleneck to much larger N. This matches the spectral picture of Section 4: norm inflation through large $\sigma_{\max}(J_x)$ selectively degrades $\kappa(H_{\text{KKT}})$ while leaving $\kappa(H_{\text{R}})$ nearly flat—the scenario in which RS is theoretically favored (Theorem 4.1).

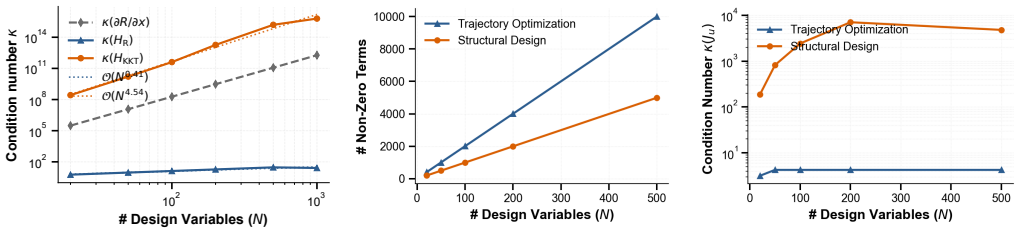


Figure 3: Left: conditioning comparison between RS and FS for the beam optimization problem. Middle: number of non-zero entries in the equality constraint Jacobian J_x for both trajectory optimization and structural design optimization. Right: condition number of J_u for FS in both problems.

Thermal topology optimization. The fourth benchmark is thermal topology optimization. We use solid isotropic material with penalization (SIMP) (Sigmund & Maute, 2013) to optimize the distribution of thermal material on a unit rod, minimizing thermal compliance $J = f^\top T$ subject to heat balance and a volume budget. With N linear bar elements the design is $\rho = (\rho_1, \ldots, \rho_N)^\top$, the interior temperatures are $T \in \mathbb{R}^{N-1}$, and the residual is $R(T, \rho) = K(\rho)T - f$:

$$\min_{\rho, T} \; f^\top T \quad \text{s.t.} \quad R(T, \rho) = 0, \;\; L_e \sum_{e=1}^{N} \rho_e = V_0, \tag{11}$$

where L_e is the element length, and V_0 the prescribed total volume. The strong form is the *second-order* conduction equation $-\frac{d}{dx}\left(\kappa \frac{dT}{dx}\right) = q$ on $(0, 1)$, with $T(0) = T_{\text{hot}}$, $T(1) = 0$, and uniform volumetric source $q \equiv 1$. Assembly gives $K_e(\rho) = \kappa_e(\rho_e)\widehat{K}$ with Laplacian template $\widehat{K}$. Thus $J_x = \partial R / \partial T = K(\rho)$ discretizes a second-order operator rather than the fourth-order bending Jacobian in the structural design optimization problem.

Table 3 summarizes RS and FS on increasing meshes. Both formulations converge on all reported cases up to $N = 500$ and agree on compliance. Unlike the beam design optimization benchmark problem, neither method fails from equilibrium violation: $\sigma_{\min}(J_x)$ remains bounded away from zero and $\sigma_{\max}(J_x)$ stays $\mathcal{O}(10^2$–$10^3)$—orders of magnitude smaller than the structural problem at comparable N. Accordingly, $\kappa(H_{\text{KKT}})$ grows with N but remains small enough that FS continues to

converge, whereas the same RS/FS split on the fourth-order beam breaks FS once $\sigma_{\max}(J_x)$ inflates. RS is consistently faster in wall time; FS requires many more iterations only at the smallest mesh ($N = 50$). This experiment therefore complements Table 2: when ill-conditioning is dominated by moderate $\sigma_{\max}(J_x)$ rather than extreme fourth-order spectral inflation, both methods remain numerically viable, even though RS is still preferable on cost.

Further discussion. Contrary to common belief, our findings reveal that matrix sparsity is not the primary determinant of algorithm performance. For instance, both trajectory optimization and structural design problems exhibit sparse constraint Jacobians $\partial g/\partial x$ with non-zero entries tightly concentrated along the diagonal. Figure 3 shows that the number of non-zero entries in J_x scales linearly with N for both problems, and we also validate the assumption that J_u is generally well-conditioned, confirming that the spectral profile of J_x is the primary governing factor. Furthermore, these two domain problems share comparable characteristics in terms of state-control dimensionality and degrees of freedom, as quantified by the ratio $\dim u/(\dim u + \dim x)$. Despite these structural and dimensional similarities, RS and FS methods perform drastically differently across the two problem classes, indicating that underlying structural features beyond sparsity and dimensionality dictate algorithmic success.

Interestingly, the structural and thermal benchmarks share nearly the same RS/FS formulation: minimize compliance subject to a discretized state equation $K(\text{design}), \text{state} = f$. Yet FS behaves very differently. It fails on fine meshes for the cantilever but converges on all reported thermal meshes. A fundamental distinction between these problems is that thermal topology is governed by a second-order differential operator

$$-\frac{\mathrm{d}}{\mathrm{d}x}\left(\kappa(x)\,\frac{\mathrm{d}T}{\mathrm{d}x}\right) = q(x),$$

whereas structural design is governed by a fourth-order operator

$$\frac{\mathrm{d}^2}{\mathrm{d}x^2}\left(E(x)\,I(x)\,\frac{\mathrm{d}^2 w}{\mathrm{d}x^2}\right) = q(x).$$

When approximating derivatives via finite differences, these operators scale with mesh size $(\Delta x)^{-2}$ and $(\Delta x)^{-4}$, respectively, where $\Delta x \sim \mathcal{O}(N^{-1})$ represents the characteristic grid spacing. Consequently, as the node count N increases, the diagonal entries of J_x in structural design grow at a rate of $\mathcal{O}(N^4)$—consistent with our experimental observation that $\kappa(H_{\mathrm{KKT}}) \sim \mathcal{O}(N^{4.54})$, leading to excessively large singular values and severe Hessian ill-conditioning.

6 Conclusion

We study the boundary between two main paradigms in constrained optimization, reduced-space and full-space methods, through the lens of equality constraint conditioning. Based on our analysis, we develop a simple criterion that provides practical guidance for selecting between reduced-space and full-space formulations based on the conditioning of the equality constraints. We validate the resulting framework across a diverse set of real-world engineering problems, where the observed solver behavior is consistent with the theoretical predictions. These results suggest that the proposed perspective is not tied to a particular application domain, solver, or problem structure, but instead captures a general numerical phenomenon arising from the formulation of equality-constrained optimization problems. More broadly, our work provides a simple, domain-agnostic, and purely optimization-based framework for understanding this problem.

While this study concentrates on how equality constraint conditioning affects optimization dynamics, inequality constraints are equally essential. When inequality constraints bound complex feasible regions, they introduce active-set transitions and non-smooth boundary conditions that substantially alter performance. Additionally, broader system properties, such as the smoothness and differentiability of the underlying model, alongside the specific search mechanisms of the numerical optimizer, exert a dominant influence on overall computational efficiency. We leave these problems for future studies.

AI USE STATEMENT

In this work, we used generative AI tools to aid or polish writing and draft sections of the paper. We have not used generative AI tools for research ideation or execution, proving mathematical claims, and generating synthetic datasets are not applicable to this work. We have reviewed all AI-assisted work, LLM-generated contents were verified and tested for correctness by all authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We provide details on the experimental setup in Appendix A. The full implementation code will be made publicly available upon acceptance.

REFERENCES

Akshay Agrawal, Shane Barratt, Stephen Boyd, Enzo Busseti, and Walaa M. Moursi. Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, 2019.

Brandon Amos and J Zico Kolter. OptNet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning (ICML)*, 2017.

Brandon Amos, Ivan Jimenez, Jacob Sacks, Byron Boots, and J. Zico Kolter. Differentiable MPC for end-to-end planning and control. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Joel A E Andersson, Joris Gillis, Greg Horn, James B Rawlings, and Moritz Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019.

Olov Andersson, Fredrik Heintz, and Patrick Doherty. Model-based reinforcement learning in continuous environments using real-time constrained optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2015.

Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Shaojie Bai, Vladlen Koltun, and Zico Kolter. Stabilizing equilibrium models by jacobian regularization. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139, pp. 554–565, 2021.

John T. Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. Society for Industrial and Applied Mathematics (SIAM), 2nd edition, 2010.

Lorenz T Biegler. An overview of simultaneous strategies for dynamic optimization. *Chemical Engineering and Processing: Process Intensification*, 46(11):1043–1053, 2007.

Lorenz T. Biegler. *Nonlinear Programming: Concepts, Algorithms, and Applications to Chemical Processes*. Society for Industrial and Applied Mathematics (SIAM), 2010.

Joshua D Deaton and Ramana V Grandhi. A survey of structural and multidisciplinary continuum topology optimization: post 2000. *Structural and multidisciplinary optimization*, 49(1):1–38, 2014.

Nikita Doikov and Yurii Nesterov. Convex optimization based on global lower second-order models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Priya L. Donti, David Rolnick, and J Zico Kolter. DC3: A learning method for optimization with hard constraints. In *International Conference on Learning Representations (ICLR)*, 2021.

Sheng Fang, Yongjin Liu, Wei Yao, Chengming Yu, and Jin Zhang. qNBO: quasi-newton meets bilevel optimization. In *The Thirteenth International Conference on Learning Representations (ICLR)*, 2025.

Ferdinando Fioretto, Pascal Van Hentenryck, Terrence W. K. Mak, Cuong Tran, Federico Baldo, and Michele Lombardi. Lagrangian duality for constrained deep learning. In *Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track*, pp. 118–135, 2021.

M. Mahdi Ghazaei Ardakani, Björn Olofsson, Anders Robertsson, and Rolf Johansson. Model predictive control for real-time point-to-point trajectory generation. *IEEE Transactions on Automation Science and Engineering*, 16(2):972–983, 2019.

Stephen Gould, Richard Hartley, and Dylan Campbell. Deep declarative networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):3988–4004, 2022.

Slavomir Hanzely, Dmitry Kamzolov, Dmitry Pasechnyuk, Alexander Gasnikov, Peter Richtárik, and Martin Takáč. A damped newton method achieves global $\mathcal o \left(\frac{1}{k^2}\right)$ and local quadratic convergence rate. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Slavomir Hanzely, Farshed Abdukhakimov, and Martin Takáč. Newton method revisited: Global convergence rates up to $o(1/k^3)$ for stepsize schedules and linesearch procedures. In *The Fourteenth International Conference on Learning Representations (ICLR)*, 2026.

Zhongkai Hao, Chengyang Ying, Hang Su, Jun Zhu, Jian Song, and Ze Cheng. Bi-level physics-informed neural networks for PDE constrained optimization using broyden's hypergradients. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.

Charles R Hargraves and Stephen W Paris. Direct trajectory optimization using nonlinear programming and collocation. *Journal of guidance, control, and dynamics*, 10(4):338–342, 1987.

Jason E. Hicken and Juan J. Alonso. Comparison of reduced- and full-space algorithms for PDE-constrained optimization. *AIAA Journal*, 51(6):1416–1426, 2013.

Ruichen Jiang and Aryan Mokhtari. Accelerated quasi-newton proximal extragradient: Faster rate for smooth convex optimization. In *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS)*, 2023.

Qiujiang Jin and Aryan Mokhtari. Exploiting local convergence of quasi-newton methods globally: Adaptive sample size approach. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

Qiujiang Jin, Ruichen Jiang, and Aryan Mokhtari. Non-asymptotic global convergence analysis of BFGS with the armijo-wolfe line search. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

James Kotary, Ferdinando Fioretto, Pascal Van Hentenryck, and Bryan Wilder. End-to-End Constrained Optimization Learning: A Survey. In *International Joint Conference on Artificial Intelligence*, pp. 4475–4482, 2021.

Dachao Lin, Haishan Ye, and Zhihua Zhang. Greedy and random quasi-newton methods with faster explicit superlinear convergence. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

Chengchang Liu and Luo Luo. Quasi-newton methods for saddle point problems. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Carlos E. Luis, Marijan Vukosavljev, and Angela P. Schoellig. Online trajectory generation with distributed model predictive control for multi-robot motion planning. *IEEE Robotics and Automation Letters*, 5(2):604–611, 2020.

Dingcheng Luo, Thomas O'Leary-Roseberry, Peng Chen, and Omar Ghattas. Efficient pde-constrained optimization under high-dimensional uncertainty using derivative-informed neural operators. *SIAM Journal on Scientific Computing*, 47(4):C899–C931, 2025.

Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, New York, NY, 2nd edition, 2006.

Elad Romanov, Fangzhao Zhang, and Mert Pilanci. Newton meets marchenko-pastur: Massively parallel second-order optimization with hessian sketching and debiasing. In *The Thirteenth International Conference on Learning Representations (ICLR)*, 2025.

Ole Sigmund and Kurt Maute. Topology optimization approaches: A comparative review. *Structural and multidisciplinary optimization*, 48(6):1031–1055, 2013.

Yongcun Song, Xiaoming Yuan, and Hangrui Yue. The admm-pinns algorithmic framework for nonsmooth pde-constrained optimization: A deep learning approach. *SIAM Journal on Scientific Computing*, 46(6):C659–C687, 2024.

Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani, Dirk Pflüger, and Mathias Niepert. Pdebench: An extensive benchmark for scientific machine learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006.

Chenrui Wang and Yixuan Qiu. The sparse-plus-low-rank quasi-newton method for entropic-regularized optimal transport. In *Forty-second International Conference on Machine Learning (ICML)*, 2025.

Sifan Wang, Ananyae Kumar bhartari, Bowen Li, and Paris Perdikaris. Gradient alignment in physics-informed neural networks: A second-order optimization perspective. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2025.

Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 2006.

A Implementation Details

Neural PDE optimization. The spatial domain $\Omega = [0, 1]$ is discretized using an N-point uniform grid $x_i = (i-1)\Delta x$ for $i = 1, \ldots, N$, with $N = 50$ and grid spacing $\Delta x = 1/(N-1)$. Boundary conditions are fixed at $v_1 = v_N = 0.1$, while interior nodes $i \in \{2, \ldots, N-1\}$ are discretized using a first-order upwind scheme for non-linear convection and a standard central difference scheme for diffusion:

$$R_i(u, S) = \begin{cases} v_i \left(\frac{v_i - v_{i-1}}{\Delta x}\right) - \nu \left(\frac{v_{i+1} - 2v_i + v_{i-1}}{\Delta x^2}\right) - u_i, & \text{if } v_i \geq 0 \\ v_i \left(\frac{v_{i+1} - v_i}{\Delta x}\right) - \nu \left(\frac{v_{i+1} - 2v_i + v_{i-1}}{\Delta x^2}\right) - u_i, & \text{if } u_i < 0 \end{cases}$$

The kinematic viscosity is set to $\nu = 10^{-3}$. Synthetic training data is generated by executing a full Newton-Raphson iteration on the discrete residual system $r(\mathbf{v}, \mathbf{u}) = \mathbf{0}$ until the residual norm falls below $\|R\|_2 < 10^{-6}$.

To train the neural operator mapping $R = R_{\boldsymbol{\theta}}$, we construct a synthetic dataset of $M = 1000$ source-state pairs $\{(x^{(k)}, u^{(k)}, v^{(k)})\}_{k=1}^{M}$, sampled from Gaussian distribution. The neural operator is trained on the PINN loss

$$L(\theta) = \sum_{k=1}^{N} |R_\theta(x^{(k)}, u^{(k)}) \frac{\partial R_\theta}{\partial x}(x^{(k)}, u^{(k)}) - \nu \frac{\partial^2 R_\theta}{\partial x^2}(x^{(k)}, u^{(k)}) - u^{(k)}|^2$$

The neural surrogate operator R_θ is formulated as a Multi-Layer Perceptron with the following structure:

- **Input Layer:** Dimension $N \in \{20, 50, 100, 200, 500\}$.
- **Hidden Layers:** Two fully connected layers of 128 neurons each, with GELU activation functions.
- **Output Layer:** Linear layer of dimension N.

The model is trained for 1000 epochs using the Adam optimizer with a constant learning rate $\eta = 10^{-3}$, achieving a final training loss below 10^{-5}. The hyperparameter $\gamma = 10^{-4}$ and the target path $\bar{u}_i = 1.5$ for $x_i < 0.5$ and 0.1 otherwise. Both RS and FS initialize control inputs at $u_t = 0$ for all t. FS initializes solution values $v_t = 0$.

Trajectory optimization. The governing ODE is given by:

$$\dot{x} = \frac{\mathrm{d}}{\mathrm{d}t} \begin{bmatrix} \beta_1 \\ \beta_2 \\ \dot{\beta}_1 \\ \dot{\beta}_2 \end{bmatrix} \tag{12}$$

where we have

$$\frac{\mathrm{d}}{\mathrm{d}t} \begin{bmatrix} \dot{\beta}_1 \\ \dot{\beta}_2 \end{bmatrix} = \begin{bmatrix} h_4 & h_5 c_{12} \\ h_5 c_{12} & h_6 \end{bmatrix}^{-1} \begin{bmatrix} h_7 s_1 - h_5 \dot{\beta}_2^2 s_{12} \\ h_5 \dot{\beta}_1^2 s_{12} + h_8 s_2 + u \end{bmatrix}$$

with parameters described as in Table 4. u is the control input. The initial state $s_0 = [0.999\pi, 0.999\pi, 0, 0]^T$ is slightly perturbed in order to avoid singularity. In the FS method, all initial states are set to s_0 at the beginning. The initial controls $u_t = 1$ for all $t = 0, ..., T-1$ for both RS and FS methods. The quadratic costs $Q = \text{diag}(1, 1, 0.1, 0.1)$ and $R = 0.1$, and the terminal cost $\Psi(x) = x^T Q x$. The forward dynamics $x_{t+1} = f(x_t, u_t)$ is computed by the fourth-order Runge-Kutta method. Both methods use exact Hessians.

Parameter	Value
l_1	$0.5L_1$
l_2	$0.5L_2$
J_1	$\frac{m_1 l_1^2}{3}$
J_2	$\frac{m_2 l_2^2}{3}$
h_4	$m_1 l_1^2 + m_2 L_1^2 + J_1$
h_5	$m_2 l_2 L_1$
h_6	$m_2 l_2^2 + J_2$
h_7	$(m_1 l_1 + m_2 L_1) g$
h_8	$m_2 l_2 g$
c_{12}	$\cos(\beta_1 - \beta_2)$
s_{12}	$\sin(\beta_1 - \beta_2)$
s_1	$\sin(\beta_1)$
s_2	$\sin(\beta_2)$
c_1	$\cos(\beta_1)$
c_2	$\cos(\beta_2)$

Table 4: Parameters in double pendulum. $m_1 = 0.2kg$ is the mass of the first rod and $m_2 = 0.2kg$ the mass of the second rod. $L_1 = 0.5m$ is length of the first rod, $L_2 = 0.5m$ is length of the second rod. $g = 9.80665$.

Structural design optimization. It considers a unit-length cantilever ($L = 1\,\mathrm{m}$) discretized with N uniform two-node Hermite cubic beam elements. The governing *fourth-order* Euler–Bernoulli equation on $(0, L)$ is

$$\frac{\mathrm{d}^2}{\mathrm{d}x^2}\left(E(x)\,I(x)\,\frac{\mathrm{d}^2 w}{\mathrm{d}x^2}\right) = q(x),$$

equivalently $(EIw'')'' = q$, where $w(x)$ is the transverse deflection and $I(x) = b\,u(x)^3/12$ for rectangular sections of fixed width b. With constant E and piecewise-constant $I_e = bu_e^3/12$ on each element,

$$E\,I(x)\,w^{(4)}(x) + 2E\,I'(x)\,w'''(x) + E\,I''(x)\,w''(x) = q(x),$$

and in the interior of element e ($I' = I'' = 0$),

$$E\,I_e\,\frac{\mathrm{d}^4 w}{\mathrm{d}x^4} = q(x).$$

We take $q \equiv 0$ on $(0, L)$; the tip load enters through the boundary conditions at $x = L$. The cantilever is clamped at $x = 0$ ($w(0) = w'(0) = 0$); a downward point load $F = -10^3\,\mathrm{N}$ is applied at the free tip $x = L$ (equivalently, $f^\top x$ in the discrete compliance objective). Each node carries transverse displacement and rotation (2 DOFs per node); the clamped root fixes those DOFs to zero. Rectangular cross-sections have fixed out-of-plane width $b = 0.01\,\mathrm{m}$; the design variables are element heights $u_e \in [h_{\min}, h_{\max}]$ with $h_{\min} = 5 \times 10^{-3}\,\mathrm{m}$ and $h_{\max} = 0.2\,\mathrm{m}$. Young's modulus is $E = 70\,\mathrm{GPa}$. Hermite cubics enforce C^1 continuity of w and discretize the bending operator; element stiffness is $K_e(u) = \frac{Eb}{12L_e^3} u_e^3 \widehat{K}$ with mesh spacing $L_e = L/N$ and height-independent template $\widehat{K}$. The volume constraint is the equality $bL_e \sum_{e=1}^{N} u_e = V_0$, where V_0 is 50% of the volume of a uniform beam at $h_{\max}$. The discrete equilibrium residual is $R(x, u) = K(u)x - f$ with $x \in \mathbb{R}^{2N}$ the free DOFs; thus $J_x = \partial R/\partial x = K(u)$ is a discrete *fourth-order* stiffness operator (two derivatives per weak-form integration by parts, twice).

Mesh sizes are $N \in \{20, 50, 100, 200, 500, 1000\}$ (Table 2). Every run starts from a uniform feasible height field at the volume budget, with $x = K(u)^{-1} f$.

Thermal topology optimization. It considers a 1D rod on $x \in [0, 1]$ with N linear bar (heat-conduction) elements. The governing *second-order* steady conduction equation is

$$-\frac{\mathrm{d}}{\mathrm{d}x}\left(\kappa(x)\,\frac{\mathrm{d}T}{\mathrm{d}x}\right) = q(x),$$

with temperature $T(x)$, conductivity $\kappa(x)$, and volumetric heat source $q(x)$. Expanding the left-hand side gives

$$-\kappa(x)\,\frac{\mathrm{d}^2T}{\mathrm{d}x^2} - \kappa'(x)\,\frac{\mathrm{d}T}{\mathrm{d}x} = q(x);$$

on each element, κ is piecewise constant ($\kappa = \kappa_e$), so

$$-\kappa_e\,\frac{\mathrm{d}^2T}{\mathrm{d}x^2} = q(x) \qquad \text{in element } e.$$

Dirichlet boundary conditions are $T(0) = T_{\mathrm{hot}} = 1$ and $T(1) = 0$; we use a uniform source $q \equiv 1$. Design variables are element densities $\rho_e \in [\rho_{\min}, 1]$ with $\rho_{\min} = 10^{-3}$. Conductivity follows SIMP Sigmund & Maute (2013) with penalty $p = 3$:

$$\kappa_e(\rho_e) = \kappa_{\min} + (\kappa_{\max} - \kappa_{\min})\rho_e^p, \quad \kappa_{\min} = 10^{-3},\ \kappa_{\max} = 1.$$

Linear bar elements use piecewise-constant κ_e on each element; assembly gives $K_e(\rho) = \kappa_e(\rho_e)\widehat{K}$ with a one-dimensional Laplacian template $\widehat{K}$. The state vector contains $N-1$ interior nodal temperatures; the discrete residual is $R(T,\rho) = K(\rho)T - f$. Hence $J_x = \partial R/\partial T = K(\rho)$ discretizes a *second-order* diffusion operator rather than the fourth-order bending Jacobian above. The volume budget matches the main text, $L_e \sum_{e=1}^{N} \rho_e = V_0$ with $V_0 = 0.5$.

Mesh sizes are $N \in \{50, 100, 200, 400, 500\}$ for RS and FS, and the results are reported in Table 3. Initial design: $\rho_e \equiv 0.5$ and $T = K(\rho)^{-1}f$. Because SIMP with $p = 3$ is non-convex, each optimization first runs a *reduced-space* warm start at $p = 1$ (200 IPOPT iterations), then continues at $p = 3$.

B An Example of Scheme-Selection Algorithm

Algorithm 1: Framework Selection Strategy Based on J_x Spectral Profile

Input : Constraint Jacobians $J_x \in \mathbb{R}^{m\times m}$, $J_u \in \mathbb{R}^{m\times n}$; Tolerances $\tau_{\text{low}}, \tau_{\text{high}} > 0$.
Output : Recommended Optimization Scheme (RS or FS).

At certain iteration, compute singular values $\sigma_{\min}(J_x)$ and $\sigma_{\max}(J_x)$;
$\kappa(J_x) \leftarrow \sigma_{\max}(J_x)/\sigma_{\min}(J_x)$;
if $\kappa(J_x)$ *is well-conditioned* ($\sigma_{\min}(J_x) \geq \tau_{low}$ ***and*** $\sigma_{\max}(J_x) \leq \tau_{high}$) **then**
    return *Both Methods*;
else
    if $\sigma_{\min}(J_x) < \tau_{low}$ ***and*** $\sigma_{\max}(J_x) \leq \tau_{high}$ **then**
        return *FS Method*;
    else
        if $\sigma_{\min}(J_x) \geq \tau_{low}$ ***and*** $\sigma_{\max}(J_x) > \tau_{high}$ **then**
            return *RS Method*;
        else
            return *Neither, apply preconditioning or constraint regularization first*;

C Proof of Theorem 4.1

The second inequality is immediate from

$$\|J_x^{-1}J_u v\| \leq \|J_x^{-1}\|\|J_u\|\|v\| = \sigma_{\max}(J_u)\,\sigma_{\max}(J_x^{-1}).$$

For the first part, suppose that $J_u v = c_1 w_1 + ... + c_m w_m$, we have

$$\begin{aligned} J_x^{-1}J_u v &= J_x^{-1}(c_1 w_1 + ... + c_m w_m) \\ &= c_1 J_x^{-1} w_1 + ... + c_m J_x^{-1} w_m \end{aligned}$$

where $c_i = w_i^T J_u v$ and hence

$$\begin{aligned}\|J_x^{-1} J_u v\| &= \sqrt{c_1^2\sigma_1^2 + ... + c_m^2\sigma_m^2} \\ &\geq \sqrt{c_{\max}^2 \sigma_{\max}^2(J_x^{-1})} \\ &= \sigma_{\max}(J_x^{-1})|c_{\max}| \\ &= \sigma_{\max}(J_x^{-1})|w_i^T J_u v| \\ &\geq c\, \sigma_{\max}(J_x^{-1})\|J_u v\| \\ &\geq c\, \sigma_{\max}(J_x^{-1})\sigma_{\min}(J_u)\end{aligned}$$

and we complete the proof.

D Proof of Theorem 4.2

Here we slightly abuse the notation for the ease of presentation. Consider the eigenvalue equation $Wx = \lambda x$ for a non-zero eigenvector $x = \begin{bmatrix} v \\ w \end{bmatrix}$ with $v \in \mathbb{R}^{n+m}$ and $w \in \mathbb{R}^m$. In block form, this yields the system:

$$Bv + J_z^T w = \lambda v, \tag{13}$$

$$J_z v = \lambda w. \tag{14}$$

We analyze two distinct cases:

Case 1: $v \in \ker(J_z)$**.** If $v \in \ker(J_z) \setminus \{0\}$, Equation 14 gives $\lambda w = 0$.

- If $\lambda \neq 0$, then $w = 0$, reducing Equation 13 to $Bv = \lambda v$.
- Pre-multiplying by v^T and normalizing $\|v\|_2 = 1$, the eigenvalue is given by the Rayleigh quotient $\lambda = v^T B v$.
- By the definition of extreme singular values of B, we have $\mu_{\min} \leq \lambda \leq \mu_{\max}$.

Case 2: $v \notin \ker(J_z)$**.** For any non-zero eigenvalue λ, Equation 14 implies $w = \frac{1}{\lambda} J_z v$. Substituting w into Equation 13 yields:

$$Bv + \frac{1}{\lambda} J_z^T J_z v = \lambda v \implies \left(\lambda B + J_z^T J_z\right) v = \lambda^2 v.$$

Taking the inner product with v (under $\|v\|_2 = 1$) gives the quadratic relationship:

$$\lambda^2 - \lambda(v^T B v) - \|J_z v\|_2^2 = 0.$$

Solving for λ via the quadratic formula yields two branches of roots:

$$\lambda = \frac{v^T B v \pm \sqrt{(v^T B v)^2 + 4\|J_z v\|_2^2}}{2}.$$

Let $\mu = v^T B v \in [\mu_{\min}, \mu_{\max}]$ and $\sigma = \|J_z v\|_2 \in [\sigma_{\min}(J_z), \sigma_{\max}(J_z)]$, and we have

Positive Branch $\lambda_+(\mu, \sigma) = \frac{\mu + \sqrt{\mu^2 + 4\sigma^2}}{2} > 0$**:** Since λ_+ is monotonically increasing with respect to both μ and σ:

$$\frac{\mu_{\min} + \sqrt{\mu_{\min}^2 + 4\sigma_{\min}(J_z)^2}}{2} \leq \lambda_+ \leq \frac{\mu_{\max} + \sqrt{\mu_{\max}^2 + 4\sigma_{\max}(J_z)^2}}{2}.$$

Negative Branch $\lambda_-(\mu,\sigma) = \frac{\mu-\sqrt{\mu^2+4\sigma^2}}{2} < 0$**:** Taking the magnitude $|\lambda_-|$:

$$|\lambda_-(\mu,\sigma)| = \frac{\sqrt{\mu^2+4\sigma^2}-\mu}{2} = \frac{2\sigma^2}{\sqrt{\mu^2+4\sigma^2}+\mu}.$$

- For a fixed μ, $|\lambda_-|$ increases with σ. Thus, evaluating at $\sigma = \sigma_{\min}(J_z)$:

$$|\lambda_-| \geq \frac{\sqrt{\mu^2+4\sigma_{\min}(J_z)^2}-\mu}{2}.$$

- As a function of μ, the expression $\frac{\sqrt{\mu^2+4\sigma^2}-\mu}{2}$ is monotonically decreasing. Hence:

$$\text{For } \mu = \mu_{\max}: \quad |\lambda_-| \geq \frac{\sqrt{\mu_{\max}^2+4\sigma_{\min}(J_z)^2}-\mu_{\max}}{2},$$

$$\text{For } \mu = \mu_{\min}: \quad |\lambda_-| \geq \frac{\sqrt{\mu_{\min}^2+4\sigma_{\min}(J_z)^2}-\mu_{\min}}{2}.$$

The maximum singular value $\sigma_{\max}(W)$ is bounded above by the maximum of λ_+:

$$\mu_{\min} \leq \sigma_{\max}(W) \leq \frac{\mu_{\max}+\sqrt{\mu_{\max}^2+4\sigma_{\max}(J_z)^2}}{2}.$$

For the minimum singular value $\sigma_{\min}(W)$:

- Globally across the spectrum, $\sigma_{\min}(W) \geq \min(\frac{\sqrt{\mu_{\max}^2+4\sigma_{\min}(J_z)^2}-\mu_{\max}}{2}, \mu_{\min})$.
- Always bounded by $\sigma_{\max}(W)$, $\sigma_{\min}(W) \leq \frac{\sqrt{\mu_{\max}^2+4\sigma_{\min}(J_z)^2}+\mu_{\max}}{2}$.

Combining these estimates into $\kappa_2(W) = \frac{\sigma_{\max}(W)}{\sigma_{\min}(W)}$ yields the lower and upper bounds:

$$\kappa_2(W) \leq \frac{\mu_{\max}+\sqrt{\mu_{\max}^2+4\sigma_{\max}(J_z)^2}}{2\cdot\min(\frac{\sqrt{\mu_{\max}^2+4\sigma_{\min}(J_z)^2}-\mu_{\max}}{2}, \mu_{\min})}$$

and we complete the proof.

E Proof of Theorem 4.3

Let v be a unit vector in $\mathbb{R}^m$. It can be uniquely decomposed into $v = v_x + v_{xc}$ where $v_x \in V_\delta$ and v_{xc} its orthogonal complement, and also $v = v_u + v_{uc}$ where $v_x \in \ker(J_u^T)$ and v_{uc} its orthogonal complement. Then we have

$$\begin{aligned}\|J_z^T v\|^2 &= \|J_x^T v\|^2 + \|J_u^T v\|^2 \\ &\geq \|J_x^T v_{xc}\|^2 + \|J_u^T v_{uc}\|^2 \\ &\geq \delta^2\|v_{xc}\|^2 + \sigma_{\min}^2(J_u)\|v_{uc}\|^2 \\ &\geq \min(\delta^2, \sigma_{\min}^2(J_u))\cdot(\|v_{xc}\|^2 + \|v_{uc}\|^2)\end{aligned}$$

On the other hand, let P_x and P_u denote the projection operator onto V_δ and $\ker(J_u^T)$, respectively. Note that

$$\|v_x\|^2 + \|v_u\|^2 = \|P_x v\|^2 + \|P_u v\|^2 = v^T(P_x+P_u)v,$$

so that it suffices to find the largest singular value of (P_x+P_u). Indeed, for any unit vector $w \in V_\delta$, we have

$$\|(P_x+P_u)w\| = \|w + P_u w\| \leq (1+\cos\theta).$$

For any unit vector w in the orthogonal complement of V_δ, we have

$$\|(P_x+P_u)w\| = \|P_u w\| \leq 1.$$

Therefore, $\sigma_{\max}(P_x+P_u) = 1+\cos\theta$ which implies

$$\|v_x\|^2 + \|v_u\|^2 \leq 1+\cos\theta,$$

and hence $\|v_{xc}\|^2 + \|v_{uc}\|^2 \geq 1-\cos\theta$ and we complete the proof.