Multi-Fidelity Robotic Co-Design via the Refinement Limit

Gordei Verbii[1], Kirill Bunin[2,3], Ilya Starikov[4]
[1]Independent Researcher [2]Kandinsky Lab (kandinskylab.ai)
[3]Moscow Independent Research Institute of Artificial Intelligence (miriai.org) [4]Moscow State University (MSU)

***Abstract*—Translating robot co-design from the lab to production demands representations that support both rapid design iteration under tight timelines *and* trustworthy high-fidelity simulation for contact-rich handling of thin, sheet-like deformables – fabric, packaging film, and bag walls. Pipelines face a dilemma: continuum simulators are accurate but too slow for design loops, while cheap discrete graph models (mass–spring, truss) may not *converge* to the physics they approximate, so a model that mis-predicts stiffness fails to transfer to the line. We supply the missing bridge. We formalize *graph-like* robots (1-skeletons of nodes and edges) as discretizations of free-form/continuum bodies at scale h, equip the hierarchy with prolongation operators, and establish: (i) geometry and the Laplace–Beltrami spectrum converge at measured rates; (ii) the discrete elastic energy Γ-converges to a continuum beam/plate limit *if and only if* refinement is all-dimensional and renormalized – otherwise the effective bending stiffness EI_h diverges or over-stiffens without bound. Where the limit exists, coarse optima prolong to near-optimal fine designs and warm-started multigrid topology optimization converges; scale-invariant controllers transfer zero-shot across resolution. These guarantees turn a naive "factor-of-two" refinement – a source of sim2real error – into a static, *simulation-side* fidelity check: a prerequisite for a multi-fidelity pipeline in which cheap discrete co-design precedes continuum-accurate validation; evidence is simulation-only, on 2-D shells.**

I. Introduction

Industrial manipulation increasingly involves *deformable* systems – soft grippers and the fabric, sheets, packaging film, and bag walls they handle, many of them mechanically *thin shells*. Moving them to production exposes a representational dilemma: a *continuum* model (FEM, material-point) is accurate but too slow for factory design loops, while a cheap *graph-like* model (mass–spring/truss 1-skeletons) is fast but need not converge to the physics it approximates. We connect the two by a single *refinement limit*: a free-form body is the $h \to 0$ limit of the graph-like body obtained by triangulating it and taking its 1-skeleton. Where this limit is well-posed, coarse and fine models are two views of one object (Fig. 1), coarse design transfers to refine, and the cheap model's fidelity can be *checked in simulation* rather than assumed.

This paper contributes foundational theory for multi-fidelity deformable-manipulation pipelines: (1) a precise statement of the robot hierarchy, its prolongation operators, and three convergence questions (geometry, energy, optimum); (2) validation that geometric and spectral structure converges at measured rates – a *fidelity benchmark*; (3) the *exact consistency conditions* under which the mechanical limit exists, with a wide-range demonstration that naive refinements have *no* limit; and (4) an account of when optimization and control transfer, with honest negatives and the design-cycle speed-ups enabled. This targets the *adaptability and repeatability* and fidelity *checking* needed for *sim2real* transfer (§IV, App. F). All numbers are from a full-resolution, *simulation-only* evaluation; we scope every mechanical claim to *2-D surfaces/shells* and make no real-robot or perception claims (gaps in Limitations).

II. Related Work

Discrete differential geometry. The cotangent Laplacian [16] and the angle-defect curvature operators [15] are the standard discrete analogues of Δ and Gaussian curvature; no discrete Laplacian satisfies all continuum properties simultaneously [20], yet mesh Laplace *spectra* provably converge to the manifold spectrum [9]. Kanai [13] established that a graph approximating a manifold is *rough- (quasi-)isometric*, not isometric, to it. **Discrete-to-continuum variational analysis.** Γ-convergence [7] is the right notion for passing from discrete energies to continua and forces *minimizers to converge to minimizers*; continuum limits of (non-convex) spring networks are classical [8, 1]. **Topology optimization** via material distribution [4, 5] and its compact SIMP codes [17, 3] underpin our structural experiments; multigrid acceleration is known to help [2]. **Differentiable soft-robot simulation and co-design** – ChainQueen [10], DiffTaichi [11], Evolution Gym [6], SoftZoo [19], and differentiable mass–spring robots [18] – and **scale-invariant/modular control** [12] motivate the transfer questions we study, complementing the morphology-as-computation view of passive dynamics [14]. These threads are exactly the ingredients a multi-fidelity pipeline for deformable manipulation must combine; we assemble them into one hierarchy and ask when the cheap end is *provably* faithful (in simulation) to the expensive one (App. C separates prior tools from what is new).

III. Proposed Method: A Refinement-Limit Framework

Notation. Throughout, $h > 0$ is the characteristic discretization length (mesh edge length, beam element size, voxel width); the continuum limit is $h \to 0$. For surface meshes we also write ε for the mean edge length ($\varepsilon \sim h$; figures use ε). A robot is denoted B_h and its continuum limit B_0.

Definition 1 (Graph-like robot)**.** $B_h = (G, \mathbf{x}, m, k, \mathcal{M})$ *where* $G = (V, E)$ *is the 1-skeleton of a simplicial discretization of a*

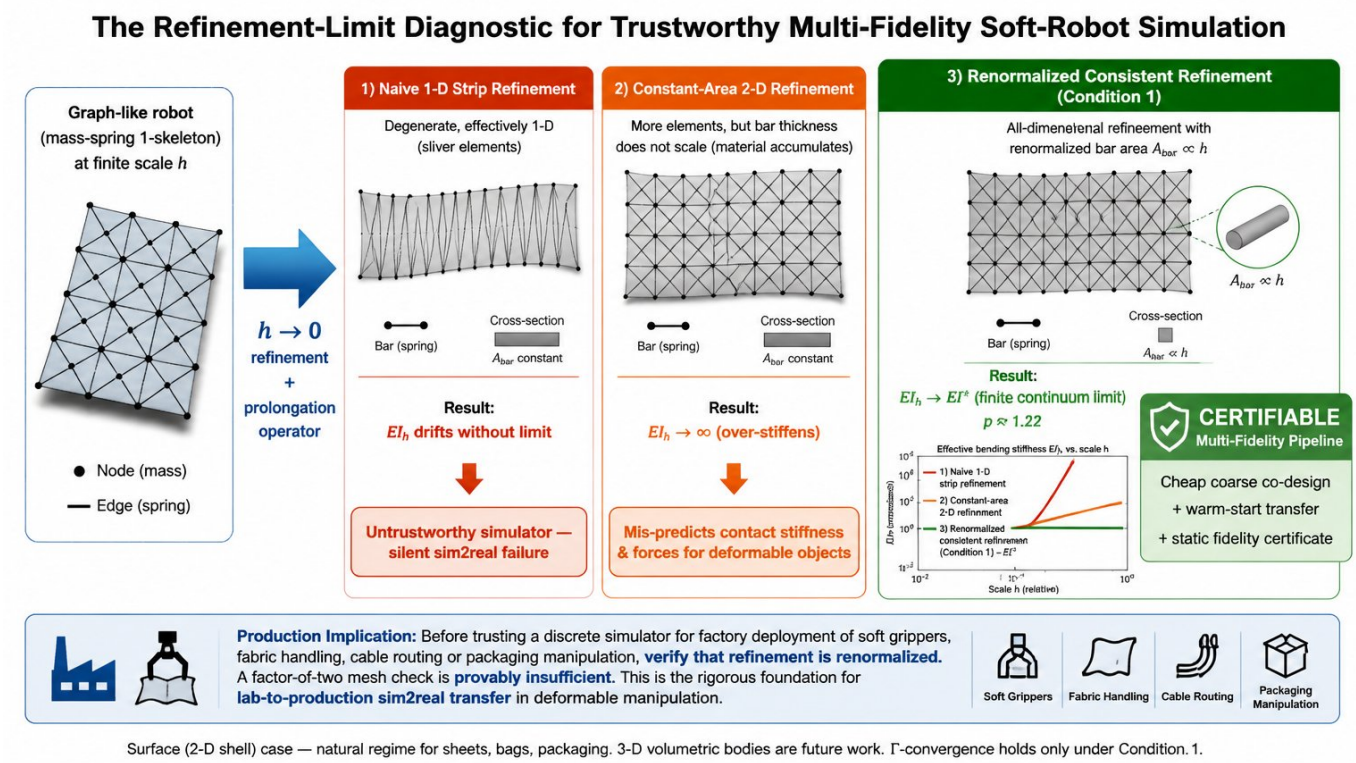

Fig. 1. **The refinement-limit diagnostic.** A graph-like (mass–spring) model of a free-form/deformable body at scale h refines toward its continuum limit three ways: the *naive strip* degenerates into slivers and EI_h drifts; *constant-area* refinement adds material and over-stiffens ($EI_h \to \infty$); only *renormalized* refinement ($A_{\rm bar} \propto h$, Condition 1) converges to a finite stiffness ($EI_h \to EI^*$, $p \approx 1.22$; cf. Fig. 2). Only the renormalized model is trustworthy, so a "factor-of-two" check is provably insufficient – the basis for a static, *simulation-side* fidelity check for a multi-fidelity pipeline on thin-shell deformables (fabric, sheets, bag walls). The analysis covers the surface (2-D shell) regime; 3-D bodies are future work.

body at scale h, $\mathbf{x} : V \to \mathbb{R}^d$ are node positions, $m : V \to \mathbb{R}_+$ masses, $k : E \to \mathbb{R}_+$ Hookean edge stiffnesses with rest lengths ℓ_e, and $\mathcal{M} \subseteq E$ the actuated (muscle) edges. A refinement *produces $B_{h'}$, $h' < h$, with a* prolongation *$P_{h\to h'}$ mapping coarse fields (designs, controls) to fine ones.*

A. Geometric layer

For a closed surface meshed by $\mathcal{T}_h$, the discrete Gaussian curvature at an interior vertex is the angle defect $K_i = 2\pi - \sum_{t\ni i} \theta_i^t$, with barycentric area weight A_i.

Proposition 1 (Topological exactness)**.** *$\sum_i K_i = 2\pi\chi(\mathcal{T}_h)$ for every h, where χ is the Euler characteristic.*

This is a combinatorial identity (sum of angle defects over a closed mesh) and is *independent of h*; it is the discrete Gauss–Bonnet theorem [15]. The pointwise density K_i/A_i converges to the smooth curvature, and the spectrum of the cotangent Laplace–Beltrami operator L_h [16] converges to that of Δ on B_0 [9]. The intrinsic graph metric is *quasi-isometric*, not isometric, to the geodesic metric [13]; the chord-to-arc distortion is $O(h^2 \sup|K|)$, so the quasi-isometry constant $L \to 1$ in Gromov–Hausdorff distance as $h \to 0$, and the finite-resolution plateau is not an asymptotic gap (App. B).

B. Physical layer and the conditions for a limit

The discrete elastic energy at displacement u is

$$E_h(u) = \sum_{e\in E} \tfrac{k_e}{2}\big(\|\mathbf{x}_i + u_i - \mathbf{x}_j - u_j\| - \ell_e\big)^2 \ (+\text{bending}). \quad (1)$$

We seek $E_h \xrightarrow{\Gamma} E_0$, a continuum hyperelastic energy. By the fundamental theorem of Γ-convergence, Γ-convergence plus equicoercivity implies $\arg\min E_h \to \arg\min E_0$ [7]: *the discrete robot's equilibria converge to the continuum robot's*. Continuum limits of spring networks of the form (1) are established under structural hypotheses on the lattice and the scaling of k_e [8, 1]. These hypotheses are not automatic; they are the crux of the framework.

Condition 1 (Consistent refinement)**.** *For E_h to admit a finite, geometry-independent Γ-limit the refinement must be (i)* all-dimensional*: the body is refined across every dimension (not only along its length), so elements do not degenerate into slivers; and (ii)* renormalized*: bar cross-sections scale as $A_{\rm bar} \propto h$ (so $k_e \propto A_{\rm bar}/\ell_e$) to keep total mass and the homogenized modulus h-independent.*

Condition 1 yields the sharp, testable prediction we exploit: define the *effective bending stiffness* $EI_h := FL^3/3c_h$ from the tip compliance c_h of a cantilever (length L, tip force F; a homogenization scaling, App. B, explains the three cases

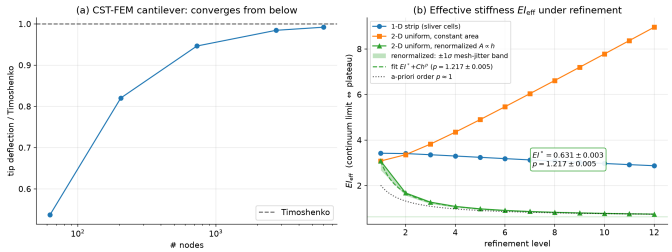

Fig. 2. **Physics convergence and the limit conditions.** (a) The constant-strain FEM cantilever converges to Timoshenko from below ($0.54 \to 0.99$). (b) Over 12 levels the naive strip *drifts* ($3.42 \to 2.87$), constant-area *over-stiffens* ($3.08 \to 8.96$), and only the renormalized lattice *plateaus* ($EI^* \approx 0.63$, $p \approx 1.22$); the $> 10\times$ spread means only renormalized refinement gives a trustworthy (sim2real-consistent) simulator.

a priori). A naive 1-D strip (refine length only) violates (i) – its sliver cells make EI_h *drift without limit*; constant-area 2-D refinement violates (ii) – it adds material and $EI_h \to \infty$; only the renormalized 2-D lattice gives $EI_h \to EI^*$ finite – the rigorous content of "a factor-of-two refinement cannot reveal a limit." In practice this is a diagnostic: before trusting a discrete simulator for a soft gripper or a grasped sheet, one must verify that refinement is renormalized, else the predicted stiffness – and contact forces – can be arbitrarily wrong.

C. Co-design and transfer-of-optimality

Co-design minimizes an objective $J_h(\rho, c)$ over a design field ρ (e.g. material density or bar areas) and a controller c, subject to the physics of B_h. The *hierarchical* scheme solves the cheap coarse problem, prolongs the minimizer, and warm-starts the fine one:

$$\rho^\star_{h_0} = \arg\min J_{h_0}; \quad \rho^{(0)}_{h_1} = P_{h_0 \to h_1} \rho^\star_{h_0}; \quad \rho^\star_{h_1} = \arg\min J_{h_1}. \tag{2}$$

Because Γ-convergence makes coarse minimizers near-optimal at the fine scale, $\rho^{(0)}_{h_1}$ lies in the correct basin and the fine solve converges in few iterations – *transfer-of-optimality*. The benefit requires a *consistent (nested) prolongation* P – a mass-preserving injection on nested grids, for which Γ-convergence places $P\rho^\star_{h_0}$ in the fine minimizer's basin (App. B): a density field on a grid admits one, but a ground-structure truss whose edge set changes under refinement does not, so its coarse optimum need not warm-start the fine one (verified in §IV). A consistent prolongation thus makes coarse-to-fine refinement *repeatable* without sacrificing converged-model accuracy.

IV. Experiments

All numbers reported here are from the full-resolution evaluation; we detail the load-bearing results below, with the hypothesis scorecard (App. Table II) and further figures in the Appendix.

Geometry and spectrum (H1–H2). Discrete Gauss–Bonnet is exact across levels, and the curvature ($O(\varepsilon^{1.21})$) and Laplace–Beltrami spectrum ($O(\varepsilon^{2.03}) \to \ell(\ell+1)$) converge at the measured rates; the quasi-isometry constant plateaus at $L \approx 1.14$, confirming quasi-isometry rather than isometry. Full numbers, the analytic torus field, and the mesh-quality (jitter) sensitivity are in App. E (Fig. 4).

The limit conditions (H3, H6). Figure 2 is the central result. The constant-strain element converges to Timoshenko (energy/Γ-convergence is confirmed separately by a Regge bending-energy test whose tip *and* stored energy reach Euler–Bernoulli, $E/E_{EB}: 1.41 \to 1.01$). The three-way EI_h contrast over 12 levels exactly matches Condition 1: a three-parameter fit $EI_h = EI^* + Ch^p$ gives $EI^* = 0.63$, $p = 1.22$ (only the constants are fitted; the limit's existence and order $p \approx 1$ follow a priori, App. B(d)), pinning a finite limit while the strip and constant-area variants diverge oppositely.

Transfer-of-optimality (H4). A five-level multigrid SIMP cascade reaches equal compliance (257.7 vs. 258.2) at $2.7\times$ fewer fine-equivalent solves, and warm-starting needs 5 vs. 23 iterations (App. Fig. 3). Optimizing the *actual 1-skeleton* (a ground-structure truss) halves compliance ($49.2 \to 24.6$) but – as predicted – does *not* warm-start (22 vs. 9 iters): its edge set is not nested. The grid/field representation supports clean transfer; the graph-native one does not (App. Fig. 8).

V. Discussion

Topological structure (Gauss–Bonnet) is exact at any resolution, and geometric, spectral, and metric structure converge ($L \to 1$ as $h \to 0$). The mechanical limit is the delicate one: it exists iff refinement is consistent (Condition 1). Where it exists, transfer-of-optimality is a property of the *representation*: a grid density field transfers, a re-meshed graph does not. Behavioral metrics poorly detect the limit, which is best *checked* statically (App. G).

Toward 3-D. The clearest way past the present 2-D scope is volumetric (tetrahedral) bodies with a mass-preserving 3-D prolongation and a Γ-limit to a 3-D hyperelastic energy; the same renormalization ($A_{\mathrm{bar}} \propto h^2$) and static check should carry over, reaching filled grippers and solid cables.

Implications for deployment. Because our evidence is *simulation-only*, we frame the EI_h trichotomy and convergence rates as a *static, simulation-side fidelity check* – an internal-consistency guarantee about the discretization, *not* a validated sim-to-real claim, and one real deployment must still pair with fabrication, closed-loop control, and hardware validation (Limitations). So read, a model is usable only where its checked quantities have converged; since behavioral metrics can miss the energy divergence, this static check — not task success alone — is the right instrument, and a consistent (grid/field) prolongation keeps coarse-to-fine refinement repeatable (App. F).

VI. Conclusion

Free-form bodies are, rigorously, the refinement limit of graph-like robots — but only under consistency conditions (Condition 1). Where they hold, geometry and spectrum converge, the elastic energy Γ-converges, and co-design and control transfer: foundational theory for multi-fidelity deformable-manipulation pipelines. The gaps to real deployment — 3-D, fabrication, perception, and hardware — are future work (Limitations).

LIMITATIONS

Assumptions. (1) *Dimension.* The quasi-isometry theorem and all experiments are for surfaces / plane bodies; 3-D filled robots are a harder, partly open regime. The surface (2-D shell) results are nonetheless the natural regime for fabric, sheet, and bag manipulation, and a stepping stone to the 3-D volumetric soft bodies that general deformable manipulation requires. (2) *Physics.* Mass–spring elasticity with penalty contact and Coulomb friction; no fabrication, actuation, or sim-to-real modeling, and no real-robot or perception experiments. This absence is a *deliberate* scope choice: the convergence framework here is a necessary *prerequisite* for – not a validation of – a trustworthy multi-fidelity pipeline. Accordingly the "fidelity certificate" throughout denotes an *internal-consistency guarantee about the discretization under our modeling assumptions*, not evidence of sim-to-real transfer; integrating it with factory-specific perception, closed-loop policies, and hardware validation is the essential next step aligned with lab-to-production goals. (3) *Consistency.* The Γ-limit requires uniform, all-dimensional, renormalized refinement; arbitrary meshing sequences need not converge. **Failure modes we observed and report.** (a) The pointwise angle-defect curvature is mesh-quality sensitive (error $\sim \delta/h^2$ under jitter), so the rate must be read from clean meshes. (b) Ground-structure truss optimization is *not* warm-startable across refinement (non-nested edge sets). (c) Locomotion distance is a weak probe of the limit; a peristaltic gait barely feels the stiffness divergence that defines it, so the rigorous probe is the static EI_h. (d) Gait fine-tuning at the finest scale gave only marginal recovery, and multi-scale GNN training did not outperform single-scale zero-shot transfer. **Future work.** 3-D tetrahedral bodies with a verified volumetric prolongation; closed-loop (per-timestep) GNN policies; provable Γ-convergence rates for the specific lattices used; fabrication-aware objectives; and hardware/perception validation to test whether the simulation-side certificate predicts real transfer.

AI USAGE STATEMENT

An AI assistant (a large language model) was used to help implement and debug the simulation code and to draft, translate and edit this manuscript. All modeling choices, mathematical statements, numerical experiments, and final claims were defined, checked, and approved by the authors, who take full responsibility for the content.

REFERENCES

[1] Roberto Alicandro and Marco Cicalese. A General Integral Representation Result for Continuum Limits of Discrete Energies with Superlinear Growth. *SIAM Journal on Mathematical Analysis*, 36(1):1–37, 2004.

[2] Oded Amir, Niels Aage, and Boyan S. Lazarov. On Multigrid-CG for Efficient Topology Optimization. *Structural and Multidisciplinary Optimization*, 49(5): 815–829, 2014.

[3] Erik Andreassen, Anders Clausen, Mattias Schevenels, Boyan S. Lazarov, and Ole Sigmund. Efficient Topology Optimization in MATLAB Using 88 Lines of Code. *Structural and Multidisciplinary Optimization*, 43(1):1–16, 2011.

[4] Martin P. Bendsøe. Optimal Shape Design as a Material Distribution Problem. *Structural Optimization*, 1(4):193–202, 1989.

[5] Martin P. Bendsøe and Ole Sigmund. *Topology Optimization: Theory, Methods, and Applications*. Springer, Berlin, 2003.

[6] Jagdeep Bhatia, Holly Jackson, Yunsheng Tian, Jie Xu, and Wojciech Matusik. Evolution Gym: A Large-Scale Benchmark for Evolving Soft Robots. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, 2021.

[7] Andrea Braides. *Gamma-Convergence for Beginners*. Oxford University Press, Oxford, 2002.

[8] Andrea Braides and Maria Stella Gelli. Continuum Limits of Discrete Systems Without Convexity Hypotheses. *Mathematics and Mechanics of Solids*, 7(1):41–66, 2002.

[9] Tamal K. Dey, Pawas Ranjan, and Yusu Wang. Convergence, Stability, and Discrete Approximation of Laplace Spectra. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 650–663, 2010.

[10] Yuanming Hu, Jiancheng Liu, Andrew Spielberg, Joshua B. Tenenbaum, William T. Freeman, Jiajun Wu, Daniela Rus, and Wojciech Matusik. ChainQueen: A Real-Time Differentiable Physical Simulator for Soft Robotics. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 6265–6271, 2019.

[11] Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. DiffTaichi: Differentiable Programming for Physical Simulation. In *8th International Conference on Learning Representations (ICLR)*, 2020.

[12] Wenlong Huang, Igor Mordatch, and Deepak Pathak. One Policy to Control Them All: Shared Modular Policies for Agent-Agnostic Control. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 4455–4464, 2020.

[13] Masahiko Kanai. Rough Isometries, and Combinatorial Approximations of Geometries of Non-Compact Riemannian Manifolds. *Journal of the Mathematical Society of Japan*, 37(3):391–413, 1985.

[14] Tad McGeer. Passive Dynamic Walking. *The International Journal of Robotics Research*, 9(2):62–82, 1990. doi: 10.1177/027836499000900206.

[15] Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H. Barr. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In Hans-Christian Hege and Konrad Polthier, editors, *Visualization and Mathematics III*, pages 35–57. Springer, Berlin, 2003. doi: 10.1007/978-3-662-05105-4_2.

[16] Ulrich Pinkall and Konrad Polthier. Computing Discrete Minimal Surfaces and Their Conjugates. *Experimental Mathematics*, 2(1):15–36, 1993.

[17] Ole Sigmund. A 99 Line Topology Optimization Code Written in Matlab. *Structural and Multidisciplinary Optimization*, 21(2):120–127, 2001.

[18] Luke Strgar, David Matthews, Tyler Hummer, and Sam Kriegman. Evolution and Learning in Differentiable Robots. In *Robotics: Science and Systems (RSS)*, 2024.

[19] Tsun-Hsuan Wang, Pingchuan Ma, Andrew E. Spielberg, Zhou Xian, Hao Zhang, Joshua B. Tenenbaum, Daniela Rus, and Chuang Gan. SoftZoo: A Soft Robot Co-Design Benchmark for Locomotion in Diverse Environments. In *International Conference on Learning Representations (ICLR)*, 2023.

[20] Max Wardetzky, Saurabh Mathur, Felix Kälberer, and Eitan Grinspun. Discrete Laplace Operators: No Free Lunch. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing (SGP)*, pages 33–37, 2007. doi: 10.2312/SGP/SGP07/033-037.

APPENDIX

A. Reproducibility and engine

The implementation details: a differentiable mass–spring engine on the 1-skeleton (semi-implicit integration with a CFL step $\Delta t = 0.2 \cdot 2\sqrt{m_i/k_i}$, penalty ground contact, anisotropic friction), discrete-geometry and FEM routines, a SIMP optimizer, and the controllers. `SMOKE` mode reproduces every code path in ~ 2 min; `FULL` reproduces all reported numbers. In particular, the notebook lets a practitioner run the *renormalized-vs-naive* fidelity diagnostic on their own lattice before committing compute to production-scale co-design.

B. Supporting analytical derivations

These derivations support the main-text claims and require no new experiments; they use only definitions and standard results from metric geometry, Γ-convergence, and asymptotic homogenization.

(a) Gromov–Hausdorff convergence of the graph metric. A map is an (L, C)-quasi-isometry if $L^{-1}d_X(x, x') - C \leq d_Y(f(x), f(x')) \leq L\, d_X(x, x') + C$. The graph metric d_G uses Euclidean chord lengths of mesh edges; d_{B_0} is the geodesic distance. For a quasi-uniform mesh of maximum edge length h and bounded aspect ratio on a compact surface of bounded curvature, an edge of arc length $s \approx h$ with local radius r has chord $\ell_e = 2r\sin(s/2r) = s\big(1 - \frac{1}{24}(s/r)^2 + O((s/r)^4)\big)$, so $|\ell_e - s|/s = O((h/r)^2)$. Summing along a geodesic of length D ($N \approx D/h$ edges) gives absolute error $\sum O(h^3/r^2) = O(Dh^2/r^2)$ and hence relative distortion

$$\left|\frac{d_G(x,y)}{d_{B_0}(x,y)} - 1\right| = O\big(h^2 \sup|K|\big) \tag{3}$$

uniformly over pairs. Therefore $d_{\mathrm{GH}}\big((V, d_G), (B_0, d_{B_0})\big) \to 0$ as $h \to 0$, and for every $\varepsilon > 0$ there is h_0 such that for $h < h_0$ the pair is $(1+\varepsilon, C_h)$-quasi-isometric with $C_h \to 0$. At a *fixed* h the chord–arc factor leaves a multiplicative term $\sim O(h^2 \sup|K|)$ bounded away from 1, visible in suprema or regression over all pairs (our $L \approx 1.14$); this is a finite-resolution phenomenon, fully compatible with metric convergence – it is *not* an asymptotic gap.

(b) The prolongation operator. Let $\mathcal{D}_h$ be the design space (piecewise-constant densities on a triangulation/grid). $P_{h\to h'} : \mathcal{D}_h \to \mathcal{D}_{h'}$ is consistent with refinement if it is linear and mass-preserving, $\int_{B_0} \rho_{h'}\, dA = \int_{B_0} \rho_h\, dA$. On nested (dyadic) grids the canonical choice is the piecewise-constant injection (each coarse cell value copied onto its 4^k fine sub-cells); it is weak-$*$ continuous and obeys $\|P_{h\to h'}\rho_h - \rho_h\|_{L^1} \leq Ch\,\|\nabla\rho_h\|$, so $P\rho_h \rightharpoonup \rho_0$ when $\rho_h \rightharpoonup \rho_0$. If the design-dependent energies satisfy $E_h(\cdot;\rho_h) \xrightarrow{\Gamma} E_0(\cdot;\rho_0)$ in a topology compatible with P (e.g. H^1 after nodal interpolation), with equicoercivity uniform on a design neighborhood, then any weak-$*$ cluster point of $P\rho_h^\star$ minimizes the limit problem; moreover, since Γ-convergence implies Kuratowski convergence of sublevel sets, for fine enough h' the point $P\rho_h^\star$ lies in the basin of a fine minimizer. A descent or optimality-criteria iteration started there converges in a number of steps essentially independent of the fine-grid size – the observed drop from 23 to 5. A ground-structure truss has design vector in $\mathbb{R}^{|E_h|}$ with $E_h \not\subset E_{h'}$ (new bars at new orientations), so no canonical linear P preserves feasibility and the quadratic-form scaling; the argument fails and warm-starting need not help (observed 9 vs. 22).

(c) Effective bending stiffness and the scaling trichotomy. For a cantilever (length L, height W, tip force F), the discrete equilibrium $u_h^\star = \arg\min_u \frac{1}{2}a_h(u,u) - F\cdot u_{\mathrm{tip}}$ has compliance $c_h = F\cdot u_h^\star$, and we define $EI_h := FL^3/3c_h$ (Euler–Bernoulli inversion; the Timoshenko inversion if shear is retained). Under Condition 1 ($k_e \sim \mathrm{const}$ because $A_{\mathrm{bar}} \propto h$ and $\ell_e \sim h$), periodic homogenization / Γ-convergence for lattice springs gives, after interpolation of nodal displacements to $v \in H^2([0,L])$,

$$\tfrac{1}{2}a_h(u_h, u_h) \xrightarrow{\Gamma} \int_0^L \tfrac{1}{2}EI\,(v'')^2\,dx + (\text{shear}), \tag{4}$$

with finite, h-independent $EI > 0$. Scaling explains the three strategies a priori. *Renormalized 2-D lattice*: number of layers $\approx W/h$, each contributing $\propto k_e \sim \mathrm{const}$; total effective width $\propto (W/h)\cdot h = W$ (fixed), hence $EI_h \to EI^* \in (0,\infty)$. *Constant-area 2-D*: A_{bar} fixed while $h \to 0$ gives $k_e \propto A_{\mathrm{bar}}/\ell_e \propto 1/h$; total width $\propto (W/h)A_{\mathrm{bar}} \to \infty$, so $EI_h \sim 1/h \to +\infty$ (over-stiffening by addition of material). *Naive 1-D strip*: a fixed number of chains has second moment of area ≈ 0 (collinear axial bars) or set by an un-renormalized per-joint torsional stiffness, so $EI_h \to 0$ or to a value disagreeing with the proper 2-D continuum limit. This is why a factor-of-two refinement cannot reveal a finite nonzero limit; the CST cantilever ($0.54 \to 0.99$) is an *independent*, high-fidelity discretization of the same continuum reference, confirming only that the target is well-posed.

(d) What is predicted a priori vs. fitted in the certificate. We make explicit which certificate numbers are derived and which are fit. The scaling argument in (c) predicts, *with no fitting*, both the *existence* and the *direction* of each case: the renormalized lattice has a finite limit $EI_h \to EI^* \in (0,\infty)$, the constant-area family diverges as $EI_h \sim 1/h \to \infty$, and

the strip has no finite non-zero limit. It also predicts the *order* of convergence: interpolating nodal displacements to H^2 makes the renormalized lattice a first-order-consistent discretization of the beam energy, so $EI_h - EI^* = O(h^p)$ with $p \approx 1$. Only three *constants* are obtained by fitting – EI^*, the amplitude C, and the exponent p in $EI_h = EI^* + Ch^p$ – via a nonlinear least-squares (Levenberg–Marquardt) regression of the 12-level renormalized sequence, i.e. a Richardson-type extrapolation of the observed plateau, with the exponent constrained only to $p \in [0.3, 4]$. The fit returns $EI^* = 0.63$ and $p = 1.22$. Two points follow. First, the fitted order $p = 1.22$ lands close to the a-priori $p \approx 1$, so the regression *confirms* the predicted rate rather than substituting for it. Second, EI^*, C, p are *problem-specific* – they depend on the chosen modulus, bar area, aspect ratio L/W, and lattice pattern – and are *not* universal constants; what the certificate asserts across problems is the *sign/existence* of the limit (predicted), with EI^* its extrapolated numerical value (measured, with the fit residual as its uncertainty).

C. Novel contributions vs. assembled prior work

Our analytic tools are standard and individually cited in §III–§IV; the contribution of this paper is their *assembly* into a single robot hierarchy and the consistency and transfer results that assembly makes provable. Table I separates the two explicitly, in response to reviewer requests to distinguish novel from assembled theory.

TABLE I
WHICH INGREDIENTS ARE PRIOR WORK AND WHICH ARE NEW HERE.

Assembled (prior work, cited)	Novel in this paper
Cotangent Laplacian and angle-defect curvature [15, 16]; Laplace-spectrum convergence [9]; quasi-isometry of the graph metric [13]; Γ-convergence and spring-network continuum limits [7, 8, 1]; SIMP topology optimization [4, 5, 17, 3] with multigrid [2]; scale-invariant / modular control [12].	(i) the graph-like-robot *hierarchy* with prolongation operators, as one object spanning geometry → physics → optimization → control; (ii) **Condition 1** (all-dimensional + renormalized) as the sharp *iff* for a mechanical limit, and the EI_h *trichotomy* read as a fidelity check; (iii) transfer-of-optimality tied to a *nested* prolongation (grid/field transfers; a re-meshed truss provably does not); (iv) the empirical finding that a *static* check dominates behavioral metrics for detecting the limit.

D. Hypothesis scorecard

TABLE II
HYPOTHESIS SCORECARD

H1 metric (quasi-isometry). GB exact; $L \approx 1.14\,(\neq 1)$. *Supported (2-D).*
H2 DDG / spectral. curvature $O(\varepsilon^{1.21})$; spectrum $O(\varepsilon^{2.03}) \to \ell(\ell+1)$. *Supported (2-D).*
H3 physics. CST→Timoshenko (0.54→0.99); bending energy → Euler–Bernoulli. *Supported (2-D).*
H4 transfer / multigrid. SIMP 2.7× fewer solves, equal compliance; warm vs. cold 5 vs. 23 iters. *Supported (2-D).*
H5 scale-invariant control. field gait transfers across 24× nodes; GNN locomotes zero-shot on an unseen resolution; shared-controller diversity → 0 vs. per-body 0.81. *Supported (2-D, with caveats).*
H6 consistency of refinement. EI_h: strip diverges, renorm → $EI^* \approx 0.63\,(p \approx 1.22)$; truss edge-set non-nested ⇒ no clean prolongation. *Conditional.*

E. Geometry and spectrum: detailed results

Geometry and spectrum (full results). Discrete Gauss–Bonnet is exact ($\sum_i K_i/4\pi = 1$ to $\sim 10^{-9}$ across 6 levels; torus total curvature → 0). On a clean mesh the curvature-density error converges at measured order $p = 1.21$ and the first six Laplacian eigenvalues at $p = 2.03$, landing on the analytic staircase $\ell(\ell+1)$ with multiplicities $1, 3, 5, \ldots$; on a torus the discrete principal curvatures converge to the exact analytic field ($K_{\max} \to 0.529$, $K_{\min} \to -1.097$). The quasi-isometry constant plateaus at $L \approx 1.14$ (mean distortion → 1.06), confirming quasi-isometry rather than isometry. Under random mesh jitter the *pointwise* curvature error grows as δ/h^2, so the clean mesh (not a jitter average) measures the true rate (Fig. 4).

F. Relevance to industrial deformable manipulation

These remarks connect the main results to the lab-to-production setting *without new experiments*; each follows from a result already established above, and all remain *simulation-side* statements pending hardware validation.

(a) The surface case is the natural regime for fabric, sheet, and bag manipulation. The quasi-isometry theorem and the geometric/spectral experiments are proved for surfaces / plane bodies. A fabric sheet, packaging film, or bag wall is mechanically a thin shell – a 2-D manifold embedded in $\mathbb{R}^3$ with negligible thickness – whose configuration space *is* a surface mesh. The geometric and spectral guarantees (exact Gauss–Bonnet, curvature $O(\varepsilon^{1.21})$, spectrum $O(\varepsilon^{2.03})$, $L \to 1$) therefore apply, unchanged, to the surface meshes used to represent such objects in a manipulation pipeline. The dimensional restriction that limits the general (3-D solid) claim is precisely satisfied here. We do not model fabric anisotropy or self-contact constitutively; that, and 3-D rods/solids for cables and filled bodies, is future work.

(b) The EI_h trichotomy is a quantified sim2real fidelity check. The effective bending stiffness $EI_h = FL^3/3c_h$ is the quantity that governs how a sheet drapes and folds and how a soft gripper responds to contact. Fig. 2(b) shows it is a function of both mesh resolution *and* the per-edge stiffness scaling: only renormalized refinement ($A_{\text{bar}} \propto h$) converges ($EI_h \to EI^*$), while constant-area refinement diverges as $1/h$ and the naive strip drifts – the three differing by $> 10\times$ at the finest level. Consequently, a pipeline that calibrates a deformable model at one resolution (e.g. for state estimation) and runs control or MPC at another will see mismatched effective stiffness – a sim2real and coarse-to-fine inconsistency – *unless* the discretization is renormalized. Because the error is not visible between adjacent levels (§III), fidelity must be checked across a wide resolution range, and the renormalized effective modulus reported and held fixed across the resolutions used at different pipeline stages. This is a necessary consistency condition, not a substitute for real-world calibration.

(c) Static checks dominate task-level metrics. Our locomotion metric is insensitive to the very stiffness divergence that defines the limit (Limitations); Γ-convergence explains why a behavioral probe can be blind to an energy divergence – the task need not excite the affected modes. The transferable prescription for evaluation under manufacturing constraints is therefore to validate deformable-object models by a static, calibrated mechanical check (a controlled bending/compliance test against a reference), not by end-to-end task success alone, which can pass while the underlying model is resolution-inconsistent and fails to transfer to a new part geometry or line speed.

(d) Consistent prolongation is a repeatability property. The density/grid representation admits a mass-preserving, nested prolongation, so coarse-to-fine refinement is predictable and the warm-start basin is guaranteed (5 vs. 23 iters); a ground-structure truss, whose edge set changes under refinement, has no canonical prolongation and the coarse optimum need not help (22 vs. 9). For a production pipeline that must re-solve grasp or placement plans as geometry varies, this is a concrete design rule: a representation with consistent prolongation gives repeatable coarse-to-fine behavior, whereas one whose connectivity is re-built each frame (e.g. a re-meshed contact graph) does not.

G. Scale-invariant control: definition and locomotion case study

A controller is *scale-invariant* if it is a function of resolution-independent body coordinates ϕ, $c = F(\phi)$, with ϕ either an intrinsic coordinate (normalized arclength) or a spectral coordinate (the Fiedler vector of L_h, which converges across scales). A message-passing graph network π_θ on G is scale-invariant *by construction*: shared weights and neighborhood aggregation define one map on graphs of any size, so $\pi_\theta \circ B_h$ is defined for all h and transfers zero-shot, giving *reusable* behaviors that survive a change of mesh resolution. The results below are a *locomotion* case study; they probe scale-invariance of the *representation* rather than an industrial manipulation task, and are reported here for completeness. A single three-parameter body-coordinate gait, optimized on a 39-node body, drives locomotion across bodies up to 949 nodes (24×); self-transfer is $\approx 100\%$ and performance degrades gradually ($\sim 18\%$/decade) while remaining positive at every scale (Fig. 5). A graph-network gait generator, scale-invariant by architecture, locomotes *zero-shot* on an unseen 441-node body (Fig. 6). A minimal model of shared- vs. per-body control reproduces a morphological-diversity collapse: shared control drives diversity to ≈ 0 while per-body control preserves it (0.81, Fig. 7) – a mechanism resonant with the precision-vs-diversity tension central to production manipulation. As noted in Limitations, gait distance is a weak probe of the mechanical limit.

H. Additional figures

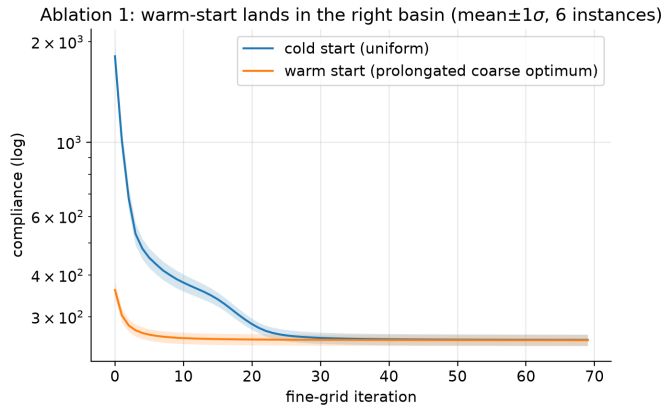

Fig. 3. **Transfer-of-optimality (H4).** On a fixed fine grid, warm-starting topology optimization from a prolongated coarse optimum reaches the solution in ~ 5 iterations versus ~ 23 cold; the multigrid cascade is $2.7\times$ cheaper overall and the margin grows with depth.

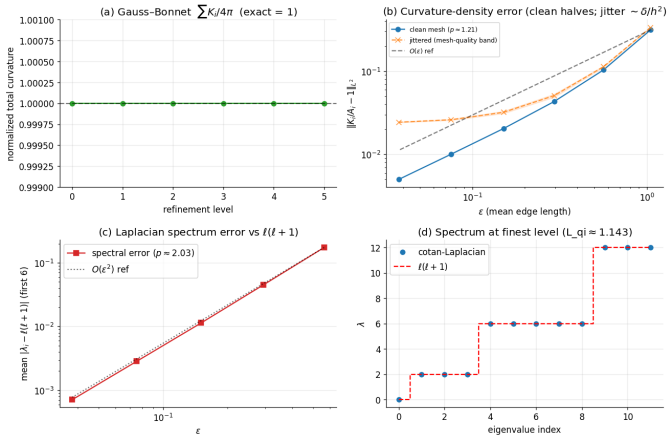

Fig. 4. Geometry/spectrum (H1–H2): exact Gauss–Bonnet; clean curvature $O(\varepsilon^{1.21})$ with a jitter (mesh-quality) band; spectrum $O(\varepsilon^{2.03})$; the $\ell(\ell+1)$ staircase at the finest level.

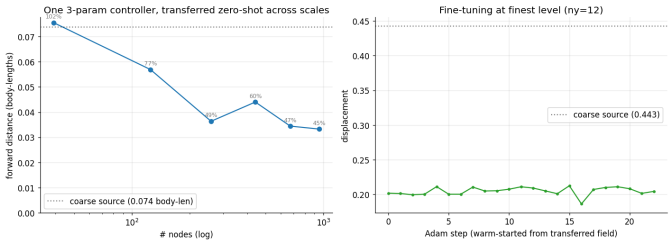

Fig. 5. Zero-shot transfer of a 3-parameter body-coordinate gait across $24\times$ nodes (self-transfer $\approx 100\%$; gradual decay), and marginal fine-tuning at the finest scale.

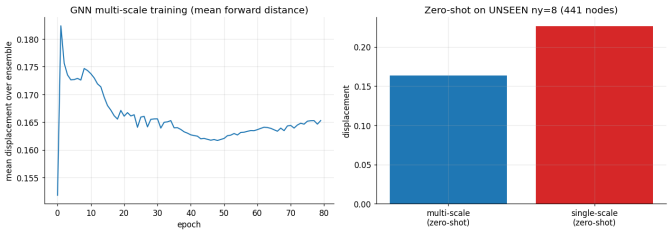

Fig. 6. A graph-network gait generator (scale-invariant by architecture) trained on a multi-resolution ensemble and evaluated zero-shot on an unseen 441-node body; both multi- and single-scale variants locomote.

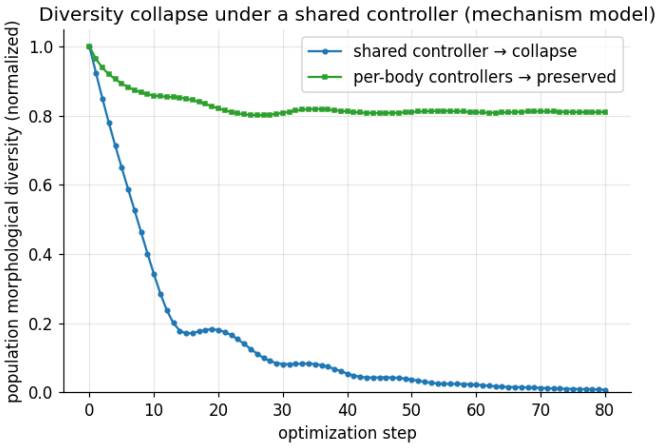

Fig. 7. Mechanism model of diversity collapse: a shared controller collapses morphological diversity to ≈ 0; per-body controllers preserve it.

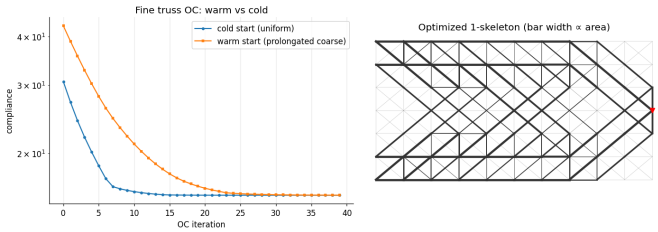

Fig. 8. Graph-native structural optimization on the 1-skeleton (ground-structure truss + Optimality Criteria): compliance halves, but the coarse optimum does not warm-start the fine truss (non-nested refinement).