TOPOLOGY-PRESERVED NEURAL SDF RECONSTRUCTION FOR SIMULATION-READY MESHING IN ADVANCED PACKAGING

Anonymous authors
Paper under double-blind review

ABSTRACT

For integrated IC to integrated system, STCO extends DTCO from chip-level to system-level co-optimization; therefore, multiphysics modeling is the core-driven force for the STCO. However, mesh generation is essential for multiphysics simulations in advanced packaging. Traditionally, these meshes require tedious manual CAD repair; it sometimes can be even slower than the simulation itself. Alternatively, a learned signed distance function (SDF) can bypass CAD repair, while reconstruct meshes directly from the easily available point scans. Whereas, it often fails to capture critical internal features, such as *vias* and narrow gaps between stacked layers. To address this issue, we propose a new CAD-preserved SDF with the following two improvements: 1) We design a new *topology*-constrained algorithm before training by pre-computing constraints from the input data. This certificate is a checkable sufficient condition that, when it holds, preserves the topology of the learned field. 2) We introduce a *discretization* certificate, a second checkable condition under which the topology is still preserved after discretization. Meanwhile, we add a *gradient-consistency* regularizer to remove stripe artifacts from flat surfaces. To be specific, we verify that every measured through-path remains open on the TSV package. The meshes generated with these improvements are then used directly for steady-state thermal simulations in a commercial multiphysics solver, without manual geometry editing. Under a shared-field vertex-placement protocol, the resulting pipeline ranks first on all six metrics for the six chip-package models, each carrying its own free-space certificate precomputed from its own scan; on three Thingi10K figurines, which have no internal corridors, the same placement rule leads on most metrics and isolates its geometric benefit from the topology claim.

1 INTRODUCTION

Multiphysics simulation helps engineers study heat flow and other effects in advanced chip packages. Such simulations need a mesh with connected material regions, open vias, and well-shaped elements (Geuzaine & Remacle, 2009; Si, 2015; Hu et al., 2018). Yet a scan does not provide this mesh directly. It may carry noise, missing surfaces, and broken connections; engineers often repair these defects before simulation (Botsch et al., 2010; Hu et al., 2020). We therefore study how to turn a package point cloud into a solver-ready mesh with little manual work.

A learned signed distance function (SDF) provides a direct route from points to a surface mesh. It represents a shape as a continuous field and smooths scan noise before extraction (Gropp et al., 2020; Ma et al., 2021; Li et al., 2024). However, a close fit to the observed surface does not ensure a correct shape inside the package. Scans may miss via interiors and narrow gaps between layers; there, training points do not say whether space is air or solid.

This missing information can change paths through a package without visibly changing surface error. In our controlled study, removing free-space protection closes four of thirteen through-paths and creates two cavities. Yet the Chamfer distance changes by less than 0.3% (Figure 3; Table 2). Thus, a surface metric can report nearly equal accuracy for two shapes with different connectivity.

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.

For a simulation mesh, that difference matters: a sealed via or gap changes the region available to the solver.

The problem begins in regions beyond surface supervision. Neural SDF losses constrain the field near observed points; Eikonal or smoothing terms prefer smooth fields elsewhere (Gropp et al., 2020; Li et al., 2024). Poisson reconstruction can also smooth over an unseen opening (Kazhdan et al., 2006; Kazhdan & Hoppe, 2013); topological losses guide training but do not protect the mesh extracted later (Hu et al., 2019; Clough et al., 2022). We call these regions the *supervision null space* (Section 3.2); we must state which of them must stay empty and check the trained field there.

We address this first problem with free-space samples computed before training. Each sample lies in a selected via or inter-layer gap and carries a conservative clearance from solid geometry; a one-sided loss discourages the SDF from labeling that sample as solid. After training, we check the field at each sample; positive values extend through a Lipschitz bound to balls of free space. A path covered by these balls stays open in the continuous field: this is our *topological certificate* (Proposition 1; Appendix E). It holds under the measured margins, the bound, and path coverage; the loss alone is not a topology guarantee.

Keeping a path open in the field is only the first step: the field must still become a mesh under an element budget. Marching Cubes yields many triangles and poor connections, and simplification can shift surfaces across narrow openings (Lorensen & Cline, 1987; Garland & Heckbert, 1997; Botsch et al., 2010). Package surfaces add a challenge: flat layers and sharp features must stay clear for meshing and simulation. Thus, a correct continuous field does not automatically give a suitable final mesh.

We address this second problem by controlling extraction and vertex motion. At a fixed vertex budget, we place vertices by curvature-adaptive sampling on the learned zero level set (Zhang et al., 2025). On the dense Marching-Cubes route, six mesh operations run under explicit displacement limits; Proposition 2 gives a sufficient condition: grid error and these displacements cannot close a protected path (Appendix F). This *discretization certificate* covers the stated Marching-Cubes chain rather than every extraction route. We also add a *gradient-consistency* regularizer (Figure 8) against flat-surface stripes from hash-encoding derivatives (Müller et al., 2022).

Together, these steps connect what must stay open in the learned field to what must stay open after meshing. Free-space protection preserves all thirteen measured paths in the controlled *CNT-TSV* comparison (Figure 3; Table 2); at a shared field and budget, our placement leads the reported metrics on six package models (Table 1). Post-processing greatly lightens *Geforce8500*, and all six package models complete the thermal-simulation workflow (Tables 10 and 12). The global condition does not hold for the tested learned-field chains; we report measured limits instead of an unconditional end-to-end certificate (Section 4.5).

The rest of the paper follows this field-to-solver path. Section 2 reviews reconstruction, topology control, and meshing; Section 3 defines the null space, the certificates, the learning objective, and the mesh routes; Section 4 tests topology, placement, mesh quality, the certificate conditions, and simulation readiness.

2 Related work

Reconstruction turns observations into a mesh under a vertex budget. Prior work addresses accuracy, topology, and usability separately; this section reviews each line and their shared gap in unobserved space.

Accuracy methods constrain the field near samples. Classical methods (Hoppe et al., 1992; Amenta et al., 2001; Carr et al., 2001) led to coordinate networks (Mescheder et al., 2019; Park et al., 2019), learned shape spaces (Chen & Zhang, 2019; Xu et al., 2019), image supervision (Wang et al., 2021; Yariv et al., 2021), local grids (Peng et al., 2020), Eikonal regularization (Gropp et al., 2020), and periodic activations (Sitzmann et al., 2020). Other methods avoid ground-truth distances or normals (Atzmon & Lipman, 2020; Erler et al., 2020; Ma et al., 2021; Peng et al., 2021; Ben-Shabat et al., 2022; Boulch & Marlet, 2022; Chen et al., 2023; Huang et al., 2023); we build on Implicit Filtering (Li et al., 2024). Hash encodings (Müller et al., 2022) speed convergence but create cell-boundary derivative jumps, targeted here by gradient consistency (Section 4.6). NeurCADRecon (Dong et al.,

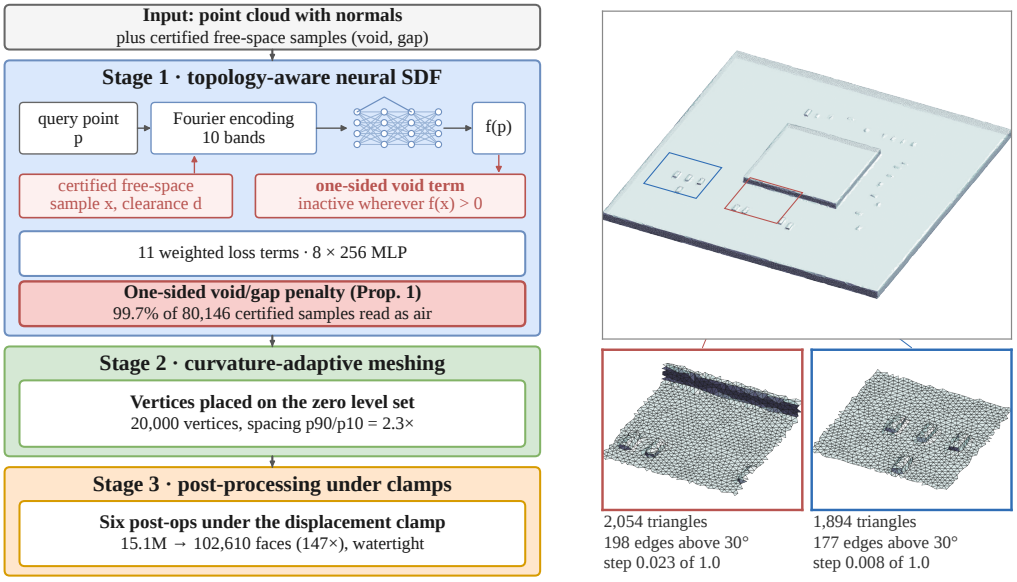

(a) The three certified stages.

(b) The board as a solver receives it.

Figure 1: **Three certified stages and a solver-ready board.** (a) Stage 1 trains a Fourier-encoded 8×256 MLP on the eleven-term objective of Table 3; Stage 2 places vertices on the zero level set and meshes them by Delaunay tetrahedralization with in/out SDF label voting; Stage 3 runs the six operations of Section 3.5 on the dense Marching-Cubes extraction under a displacement clamp, each stage box printing one quantity measured on its own artifact; Stage 2's budget-exact Delaunay mesh is scored in Table 1. (b) *Geforce8500* after post-processing, $147\times$ lighter than the 15.1M-face extraction. Zoom windows sit where sharp-edge density concentrates, not by hand (red: layer step; blue: flat region).

2024) regularizes Gaussian curvature, not through-hole topology. These terms leave hidden free space unconstrained.

Persistent-homology losses constrain Betti numbers (Hu et al., 2019; Clough et al., 2022; Gabrielsson et al., 2020; Edelsbrunner & Harer, 2010). They differentiate through a persistence diagram at every step. Such a training penalty, however, gives no guarantee once the field is extracted and post-processed. Our certificates are local, geometric, and computed before training: each sample carries a clearance the field may not violate in one forward evaluation. Proposition 2 then carries the per-via guarantee of Proposition 1 to the output mesh. We know of no prior certificate for the null space of point-cloud supervision.

Meshing starts with extraction. Edge collapses (Hoppe, 1996), planar proxies (Cohen-Steiner et al., 2004), and feature-sensitive contouring (Kobbelt et al., 2001) control count and sharp edges. Learned extractors, including Neural Marching Cubes (Chen & Zhang, 2021), Neural Dual Contouring (Chen et al., 2022; Ju et al., 2002), and VoroMesh (Maruani et al., 2023), use low-resolution grids. DMTet (Shen et al., 2021) and FlexiCubes (Shen et al., 2023) relax cells but keep a fixed layout, whereas PoNQ (Maruani et al., 2024) abandons it. LMR (Zhang et al., 2025) is closest to our design; we add topology-aware supervision and simulation post-processing. Next, QEM simplification (Garland & Heckbert, 1997) and feature-preserving remeshing (Botsch et al., 2010) repair surfaces. Finally, TetGen (Si, 2015), Gmsh (Geuzaine & Remacle, 2009), and TetWild (Hu et al., 2018; 2020) improve volumetric connectivity; primitive recovery remains separate (Schnabel et al., 2007). These tools do not recover sealed vias or enforce semantics. Our post-processing combines watertightness, aspect ratio, and explicit primitives under certificate-chain displacement clamps.

No prior line gives one checkable statement for accuracy, topology, and usability on the same field and budget: Proposition 1 protects the field, Proposition 2 carries it to the mesh, and the remaining properties are measured on shipped artifacts (Section 4.5).

3 METHOD

3.1 OVERVIEW

The input is a point cloud $P = \{p_i\}_{i=1}^{N}$ normalized to $[-0.5, 0.5]^3$. We learn a neural SDF $f_\theta : \mathbb{R}^3 \to \mathbb{R}$ (Section 3.3). Two routes mesh this field: curvature-adaptive Delaunay meshing at an exact vertex budget (Section 3.4), or dense Marching Cubes with feature-preserving post-processing (Section 3.5). Each stage solves a different failure mode (Section 3.2).

3.2 THE SUPERVISION NULL SPACE

Any loss fitting a neural SDF to a point cloud acts only near the observations. Let $\mathcal{N}_\tau(P) = \{x : \mathrm{dist}(x, P) \le \tau\}$ be the τ-neighborhood of the surface; every data term in Section 3.3 satisfies $\mathrm{supp}(\mathcal{L}_{\mathrm{data}}) \subseteq \mathcal{N}_\tau(P)$. We call the complement $\Omega_0 = [-0.5, 0.5]^3 \setminus \mathcal{N}_\tau(P)$ the *supervision null space*. Data-term minimizers agree on samples yet may disagree arbitrarily on Ω_0. Engineering geometry must be correct there. This disagreement has three forms. First, *topology* decides whether an unsampled deep via stays open or fills. Next, *field regularity* controls derivative jumps at hash-cell boundaries, visible as stripes on planar faces. Finally, *discretization* of the zero level set decides the mesh and post-processing motion. Denser sampling of the same visible surfaces cannot close these gaps. We therefore constrain Ω_0 with two checkable certificates and one measured regularizer, one per failure mode (equation 1, the displacement clamps of Section 3.5, and equation 2); Appendix K details the construction and its distinction from a prior.

3.3 TOPOLOGY-AWARE SDF LEARNING

Network. f_θ is a Fourier-encoded 8-layer, 256-unit MLP (Tancik et al., 2020). An optional hash encoding (Müller et al., 2022) adds detail. This encoder can disturb the early field, so it warms up late at $0.1\times$ learning rate (Appendix A).

Edge-preserving bilateral filtering. Following Li et al. (2024), each query first moves onto the zero level set along the gradient, and is then filtered against its nearest surface points, with bandwidth $\sigma = 2\,\bar{d}_{\mathrm{1nn}}$ and normal width 15° (Appendix A).

Void/gap topology protection. Through-hole and gap interiors are never sampled, so f_θ can seal them. We pre-compute free-space samples x_j inside them. Each carries a conservative clearance $d_j > 0$ to the nearest *solid*; it lies in certified air and must stay positive. We penalize shortfall below this clearance:

$$\mathcal{L}_{\mathrm{void}} = \lambda_{\mathrm{void}} \Big[\operatorname*{mean}_j \mathrm{ReLU}\big(d_j - f_\theta(x_j)\big)^2 \,(1 + w_j^d) \;+\; \tfrac{1}{2} \operatorname*{mean}_j \big|f_\theta(x_j) - d_j\big| \Big]. \tag{1}$$

Our quadratic term is a *one-sided fill penalty*, active only at samples with $f_\theta(x_j) < d_j$. This depth weight $w_j^d = d_j / \max_k d_k$ favors deep-hole samples, sealed first. An auxiliary L_1 term pulls $f_\theta(x_j)$ toward d_j in air. Voids and gaps use separate batches (8,192 and 4,096 pts per step; Appendix A). Points alone do not protect a via; the Lipschitz bound grows each point into a safe ball. Proposition 1 (Appendix E) formalizes this guarantee.

Proposition 1 (through-hole preservation)**.** *Let $\{(x_j, d_j)\}_{j=1}^{m}$ with $d_j > 0$ be the certified free-space samples of equation 1, with the trained field attaining margins $f_\theta(x_j) \ge m_j > 0$, and assume f_θ is L-Lipschitz. Then $f_\theta > 0$ on the open ball $B(x_j, m_j/L)$ for every j; consequently any curve γ traversing a via covered by these balls (consecutive certified samples spaced closer than $\min_j m_j/L$ suffice) satisfies $\gamma \cap \{f_\theta \le 0\} = \emptyset$, i.e. the through-hole remains open in the reconstruction.*

Sign and far-field constraints. Three one-sided terms keep spurious negative pockets out of unobserved regions: the field must stay positive beyond 5σ from P, on an air layer beyond 3σ, and on the bounding-box faces and corners (formulas and thresholds in Appendix A). The Eikonal regularizer (Gropp et al., 2020) and a finite-difference total-variation term complete the objective.

Multi-scale gradient consistency. Hash encodings create derivative jumps at cell boundaries, visible as stripes on planar faces. We therefore compare derivatives across boundaries. For random

points x we draw Gaussian offsets $\delta \sim \mathcal{N}(0, \sigma_\ell^2 I)$, with σ_ℓ at $1.5\times$ the cell size of a random coarse level ($\ell \in \{1, 2, 3\}$). This offset usually places the pair $(x,\ x{+}\delta)$ across one cell boundary, and we penalize

$$\mathcal{L}_{\mathrm{gc}} = \lambda_{gc}\,\mathrm{mean}\big(\|\nabla f_\theta(x) - \nabla f_\theta(x{+}\delta)\|_2^2\big). \tag{2}$$

Training strategy. Adam trains each scene under warm-up and cosine annealing, with gradient accumulation on large clouds. Two data-side corrections matter on packages: rescaled sampling noise on elongated boxes, and up-weighting for sparse vertical walls. Loss weights are in Table 3 and the schedule in Appendix A.

3.4 Curvature-adaptive vertex generation and Delaunay meshing

We place M vertices on the zero level set (Zhang et al., 2025) at an exact budget, instead of contouring a dense grid. First, candidates move onto it by SDF gradient steps. Next, mean curvature guides resampling toward creases, corners, and small bores. Then, repulsion relaxes spacing, and upsampling fills under-resolved regions. Finally, Delaunay tetrahedralization meshes the vertices, and a voting rule over SDF values and gradients labels each tetrahedron inside or outside (Appendix A). The interface is watertight and orientable. It is not bound to a fixed voxel layout; certified topology still needs Proposition 2.

3.5 Feature-preserving lightweighting and simulation-oriented post-processing

Post-processing stages. Six stages run in order on the dense Marching-Cubes extraction, each under an explicit displacement clamp (parameters in Appendix A). The extraction is too dense, so *vertex clustering* collapses it on a uniform grid to $\sim$100K. *Planar sharpening* quantizes normals to 18 directions. These labels mark each vertex as plane interior, edge, or corner, sending it onto the matching plane or intersection under a 3%-diagonal clamp. *Distance-aware selective QEM* decimates (Garland & Heckbert, 1997) under per-region quotas, protecting creases; a binary search hits the budget exactly. *Multi-domain decomposition* splits the mesh into the input model's part domains, each remeshed under its own budget. *Analytic feature snapping* snaps vertices onto planes, vias, and creases within 2% of the diagonal. *Aspect-ratio repair* targets $\mathrm{AR} < 5$ and validates watertightness, winding, non-manifold edges, degenerate faces, and Euler characteristic.

Certificate chain under discretization. Each stage can undo a field-level certificate, so each carries a displacement bound. In simple terms, the mesh stays safe while vertex motion plus grid error stays below the smallest field margin. Proposition 2 formalizes this idea: the certificate of Proposition 1 survives Marching Cubes and all six stages (Appendix F). The condition is checkable on shipped artifacts: m_j, h, and $\sum_k \varepsilon_k$ are measured exactly, while L is estimated empirically (Section 4.5).

Proposition 2 (certificate chain under discretization and post-processing)**.** *Let f_θ be L-Lipschitz, let $\{(x_j, d_j)\}_{j=1}^m$ be certified free-space samples, with the trained field attaining margins $f_\theta(x_j) \geq m_j > 0$, and put $R_j = m_j/L$. Let the zero level set be extracted by Marching Cubes on a grid of edge length h, and let the stages $k = 1, \ldots, K$ above displace every vertex by at most ε_k. If*

$$\sum_{k=1}^{K} \varepsilon_k \;+\; \sqrt{3}\,h \;<\; \min_j R_j, \tag{3}$$

then the output mesh does not intersect the ball $B(x_j, \rho_j)$ with $\rho_j = R_j - \sqrt{3}h - \sum_k \varepsilon_k > 0$, for every j. Consequently any curve γ traversing a via covered by the balls $\{B(x_j, \rho_j)\}$ (consecutive certified samples spaced closer than $\min_j \rho_j$ suffice) stays in a single component of the complement of the output mesh, i.e. the certified tunnels cannot be closed by discretization or by post-processing.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Datasets. Nine point clouds enter the evaluation. *CNT-TSV* (200K points) stacks shells over CNT bundles and through-silicon vias; its corridor network contains 126,174 certified free-space samples. *PackagedChip* (2M) is a full assembly with a 512^3 reference of 57 bodies and 15.7M faces; its voids and inter-layer gaps contain 143,089 certified samples. *Geforce8500* (800K) and *780 Chipset* (500K) test thin layers: their shortest bounding-box axes are 5.5% and 2.4% of the longest. Thus, a grid resolving the footprint barely resolves the stack. *Through-Hole* (200K) is a small single body, whereas *Serdes* (400K) scans a production-package GND connector ribbon. Finally, three Thingi10K figurines (Zhou & Jacobson, 2016) (*Thingi-44234*, *-47984*, *-354371*; 99,960 points each) test placement on curved surfaces without corridors (Appendix D).

Metrics. Surface accuracy uses CD, NC (k=5), F1 at 0.008, ECD for sharp edges (Chen & Zhang, 2021), and VSD, following Zhang et al. (2025). Topology uses cubical-complex Betti numbers on raw 512^3 MC occupancy (Edelsbrunner & Harer, 2010): b_1 counts tunnels and b_2 sealed cavities (Appendix A). The field check uses sealed fraction and clearance error on the certified samples of equation 1. Mesh quality uses watertightness, non-manifold edges, minimum angle, aspect ratio, and genus.

Baselines. We first evaluate every method on the *same frozen field* (Table 1), isolating vertex placement. NMC (Chen & Zhang, 2021), NDC (Chen et al., 2022), and VoroMesh (Maruani et al., 2023) are surrogates, namely placement rules without learned components. In contrast, MC uses the requested budget and SPR (Kazhdan & Hoppe, 2013) fits its own screened-Poisson field. All entries read the same gradient normals; SPR also gets oracle depth selection (Appendix C). NeurCADRecon is omitted because its official implementation was unavailable in our environment (Appendix D).

Implementation and cost. We use PyTorch (Paszke et al., 2019) on V100-32GB. Each scene trains at LR 10^{-3} for 70K–120K steps; batch sizes and wall-clock appear in Appendix A. Certificate samples are pre-computed before training in about one GPU minute; after training, checking the field at every sample and Delaunay meshing each take one GPU pass, and post-processing two CPU minutes (Table 12).

4.2 VERTEX PLACEMENT UNDER A SHARED FIELD

Under the shared-field protocol, ours ranks first on all 36 cells on the six package models (VSD_{avg}/$\text{VSD}_{\max}$/CD/ECD in 10^{-4}/10^{-3}/10^{-5}/10^{-3}; best in bold; Table 1). It ties SPR only on Through-Hole F1 of 1.000.

This overall lead is largest where topology and thin features stress voxel-bound rules. *Geforce8500* ECD is $5.7\times$ below NDC, and *780 Chipset* VSD_{avg} is $37\times$ below NDC while preserving thin layers. *CNT-TSV* ECD is $1.6\times$ below the next-best; *PackagedChip* and *Serdes* lead all six cells. Figure 2 shows the $37\times$ gap: competing rules snap to a grid or scatter off surface, whereas ours traces the slot. The same protocol ranks ours first on all 18 cells for three Thingi10K figurines (Table 4).

The reference is our field's zero surface, not the original CAD, and the three surrogates lower-bound their published systems (Appendix C). Our 20K budget is exact on every cloud, whereas grid-bound rules return what their grids allow. On repaired Thingi fields (Appendix D), ours leads 15 of 18 cells. This is the Stage 2 evaluation; the production route of Section 3.5 reads the same field.

End-to-end reconstruction. Without a frozen field, every method builds from the same cloud and is scored against ground-truth CAD (Table 5, Appendix D). Ours beats SPR normal consistency on all six packages. On *CNT-TSV*, SPR seals the corridors: its CD is $33\times$ ours and F1 is 0.768 versus 0.887.

NDC has lower curvature error at 4–$104\times$ fewer elements. However, its grid cannot represent the certified through-holes, and none of its nine outputs is watertight or manifold. On *CNT-TSV*, SPR edge distance beats our raw extraction.

Table 1: **Vertex placement on a shared SDF field.** Six chip-package models, six placement rules, 36 model–metric cells, all at the same requested budget of 20K vertices. SPR is one realization of a non-bit-reproducible solver; run-to-run spread $1.5/14/0.6/0.38/2.5/3.9\%$ per metric.

Dataset	Method	$\text{VSD}_{\text{avg}}\downarrow$	$\text{VSD}_{\max}\downarrow$	CD↓	NC↑	F1↑	ECD↓
Through-Hole	NMC	112.37	53.91	23.65	0.959	0.202	0.044
	NDC	4.38	4.93	3.36	0.973	0.717	0.038
	VoroMesh	179.10	104.94	57.94	0.993	0.288	0.024
	MC (equal budget)	0.95	2.49	0.93	0.994	0.991	0.009
	SPR	1.24	1.86	0.59	0.996	**1.000**	0.004
	Ours	**0.42**	**0.99**	**0.45**	**0.996**	**1.000**	**0.002**
Geforce8500	NMC	34.46	50.91	12.80	0.849	0.292	0.091
	NDC	4.67	5.27	7.10	0.912	0.360	0.046
	VoroMesh	217.64	134.86	95.22	0.823	0.110	0.068
	MC (equal budget)	1.01	3.24	3.24	0.997	0.676	0.020
	SPR	0.91	2.95	2.62	0.997	0.814	0.018
	Ours	**0.22**	**1.53**	**1.97**	**0.998**	**0.880**	**0.008**
780 Chipset	NMC	22.46	4.48	8.19	0.880	0.311	0.087
	NDC	5.84	5.79	7.25	0.903	0.348	0.046
	VoroMesh	213.10	100.46	91.88	0.907	0.150	0.055
	MC (equal budget)	3.41	2.87	2.87	0.998	0.748	0.015
	SPR	1.36	3.02	2.51	0.997	0.831	0.018
	Ours	**0.16**	**0.94**	**1.97**	**0.998**	**0.862**	**0.004**
PackagedChip	NMC	18.33	10.56	340.94	0.541	0.222	0.145
	NDC	11.45	6.67	10.82	0.738	0.262	0.084
	VoroMesh	211.91	75.38	115.53	0.435	0.029	0.514
	MC (equal budget)	22.09	12.52	21.54	0.625	0.200	0.152
	SPR	195.10	121.35	197.18	0.356	0.073	1.523
	Ours	**1.58**	**3.19**	**7.83**	**0.794**	**0.358**	**0.056**
CNT-TSV	NMC	53.81	12.18	9.63	0.933	0.317	0.060
	NDC	4.32	4.17	4.82	0.949	0.531	0.037
	VoroMesh	225.81	104.97	103.53	0.854	0.146	0.050
	MC (equal budget)	1.79	3.67	2.23	0.994	0.839	0.014
	SPR	8.38	23.41	5.06	0.950	0.870	0.015
	Ours	**0.36**	**0.93**	**1.89**	**0.994**	**0.905**	**0.009**
Serdes	NMC	20.57	77.37	9.74	0.856	0.303	0.075
	NDC	6.63	8.75	8.35	0.874	0.399	0.085
	VoroMesh	181.63	127.30	141.43	0.749	0.203	1.079
	MC (equal budget)	6.71	10.89	8.10	0.869	0.395	0.085
	SPR	28.20	41.41	1893.49	0.721	0.496	36.128
	Ours	**0.42**	**3.42**	**2.63**	**0.958**	**0.780**	**0.020**

4.3 Adaptivity and lightweighting

Curvature adaptivity. On *Geforce8500*, density in the high-curvature band is $3.91\times$ that of the low band (70th vs. 30th percentile), versus $1.44\times$ without curvature (Figure 6). This shift explains the ECD margins in Table 1.

Lightweighting. The 2048^3 MC output becomes watertight and manifold (Table 10, Appendix H). Near-degenerate faces below 10° drop from 44,270 to 7, the mean minimum angle rises to 52.1°, and faces fall from 15.1M to 102,610 ($147\times$).

4.4 Simulation readiness on real packages

CNT-TSV decomposes into 14 watertight domains; its final ≈20K-face mesh has 1.26×10^{-5} Chamfer from the input. *PackagedChip* snaps 591 planes, 96 cylinders, and 103 creases onto explicit primitives. The result is watertight, has zero non-manifold edges and $\text{AR}_{\max} = 4.9996 < 5$, and passes the Euler/genus check (Figure 5; Table 12; Appendices H and B).

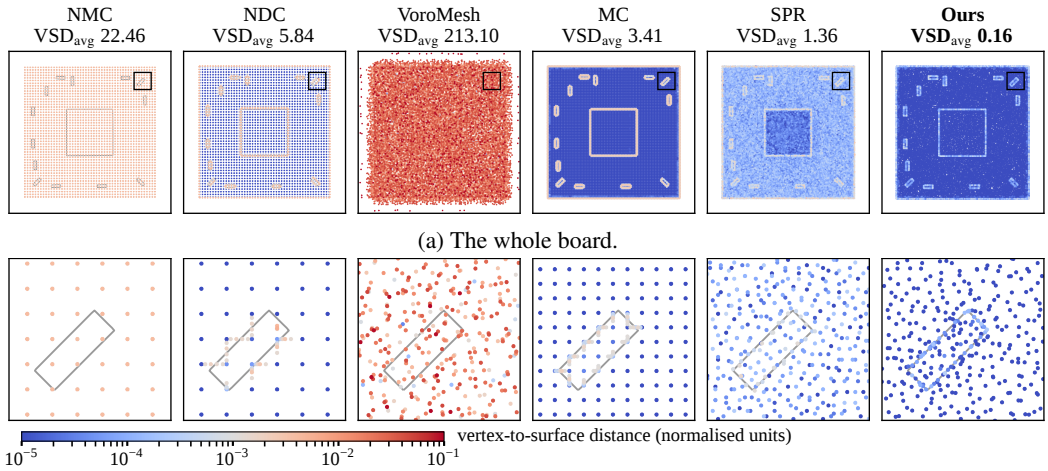

(a) The whole board.

(b) The rim slot boxed in (a).

Figure 2: **Where the vertices land, at one shared field and one budget.** *780 Chipset* at a requested budget of 20K vertices, coloured by unsigned distance to the 512^3 zero surface on a log scale; the grey outline is the mid-plane section of the reference mesh, common to all six panels. The NMC and NDC surrogates inherit their coverage from the same 512^3 grid—cell centres versus edge crossings—so the two panels look alike; they differ in landing error: 74.9% of NMC's vertices lie beyond 10^{-3}, against 21.0% of NDC's and 96.2% of VoroMesh's, while ours holds every vertex inside 9.4×10^{-4} and traces the slot outline. The last column looks dense because its points land on the outline, not because it holds more points: all six rules answer the same 20K request, so the density adds no meshing cost (Section 4.1).

All six package meshes import into a multiphysics simulator and solve stationary heat conduction without user-guided geometry editing. *Geforce8500* takes 1.0 s at 28,953 tetrahedra after the solver's default edge repair; *Serdes* takes 54 s at 1.73M elements (Appendix I).

4.5 Verifying the certificate chain

Three *CNT-TSV* chains test equation 3 (Appendix G): A (field $\rightarrow$ MC^{512}) tests discretization, B (CAD $\rightarrow$ mesh) post-processing, and C is the negative control. B holds: $7.9\times$ margin, all 126,174 samples satisfy equation 3, and Betti is unchanged. It is sufficient, not necessary; failure need not imply topology loss. It fails on both learned-field chains (0.47% of A samples have $m_j \leq 0$), yet holds for 84.1% at empirical L_{p99} and A reproduces $b_1 = 13$. Chain C has higher $\sum_k \varepsilon_k$; 13.4% of its balls intersect.

4.6 Ablation study

Variant A (voxel vertices) degrades every metric by up to two orders of magnitude (Table 8, Appendix H). Variant B (no curvature) trades crease for uniform accuracy: ECD worsens $6.6\times$ on crease-rich *Geforce8500*, whereas smooth CD/NC improve slightly. Simulation needs the creases and bores carrying boundary conditions.

Surface metrics miss topology. Two controlled 20K-step *CNT-TSV* trainings differ only in $\mathcal{L}_{\text{void}}/\mathcal{L}_{\text{gap}}$ (Figure 3, Table 2). At identical surface accuracy (CD 3.378 vs. 3.367, $\times 10^{-3}$), removing the loss seals two cavities (b_2: $0 \rightarrow 2$) and pinches four tunnels ($13 \rightarrow 9$). Sealed fraction rises $2.7\times$ in voids and $4.0\times$ in gaps; clearance error rises 2.4–$3.1\times$, and $\min_j f_\theta$ changes from -0.014 to -0.033. Both runs keep GT-SDF supervision, so residual sealing shows what it cannot prevent alone.

The 2×2 study (Table 14, Figure 8, Appendix J) tests encoder decay and $\mathcal{L}_{\text{gc}}$. Either alone holds planar deviation near $0.36°$. However, with both off, deviation ($0.74°$), stripe spectrum ($1.7\times$), and ECD ($+59\%$) degrade together. The shipped pair is not the table optimum because neither single-mechanism arm was validated beyond this model.

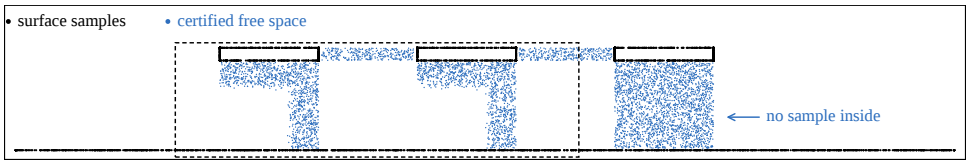

(a) The input cloud and the certified free-space samples on the cut plane; the dashed box is the window of (b) and (c).

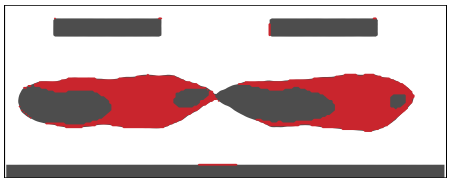

(b) Certificate off: the corridors fill in.

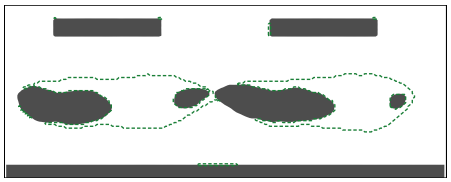

(c) Certificate on (ours): the corridors stay open.

Figure 3: **A sealed via is invisible to surface metrics.** *CNT-TSV*, one plane cut ($y = +0.137$) through the 512^3 MC occupancy of two fields differing only in the free-space certificate of equation 1. (a) the input cloud never samples a via interior, so the 126,174 clearance samples computed before training carry the only information about it. (b) without the certificate 84,479 certified voxels turn solid, b_1 falls from 13 to 9, and two cavities appear; (c) with it they stay air. Chamfer moves by 0.3% between the two fields, so Proposition 1 constrains the field, not the metric.

Table 2: Ablation of the void/gap protection loss on *CNT-TSV*: two controlled 20K-step runs differing only in $\mathcal{L}_{\text{void}}$ (GT-SDF active in both). Sealed fraction counts certified free-space samples with $f_\theta \leq 0$; clearance error and CD are in 10^{-3} normalized units.

Variant	$b_1\uparrow$	$b_2\downarrow$	**Sealed frac.** $\downarrow$ **(%)**		**Clearance err.** $\downarrow$		**CD**$\downarrow$
			void	gap	void	gap	
with $\mathcal{L}_{\text{void}}$	**13**	**0**	**0.28**	**0.81**	**0.53**	**0.85**	3.378
w/o $\mathcal{L}_{\text{void}}$	9	2	0.77	3.24	1.27	2.63	**3.367**

Cross-model without GT-SDF. A user without a CAD reference has no GT-SDF supervision, so we repeat the pair with it disabled (Table 9, Appendix H). Without protection, *CNT-TSV* seals 16.0% of its gap samples by 20K steps and inflates the clearance error by 43–96×; *PackagedChip* peaks at 72.8%. With protection, *CNT-TSV* is clean from 10K steps and *PackagedChip* stays below 10.4%; the protection is not free, raising CD on *PackagedChip* from 12.5 to 18.6×10^{-3}.

Sensitivity. Around the sign-supervision operating point, fidelity and topology hold to within 1% CD and to an invariant genus, while ECD swings 6.4× on a nearly edge-free model (Table 11, Appendix H). **Not ablated**: flat-floor backfill and the 20K budget, each explained in Appendix H.

5 Conclusion

We present a pipeline from chip-package point clouds to lightweight meshes for multiphysics simulation, without user-guided geometry editing. First, a topological certificate protects unobserved vias and inter-layer gaps during SDF learning. Next, a discretization certificate bounds extraction and post-processing, while gradient consistency reduces planar striping. Experiments preserve through-paths on *CNT-TSV*, give accurate surface discretization on six package models, and improve mesh quality across package and Thingi10K models. All six package meshes complete automatic import, volume meshing, and steady-state heat-conduction simulation. Future work will expand certificate coverage and package-family generalization; Appendix L states the limitations.

AI USE STATEMENT

Generative AI tools assisted in organizing related literature. The authors reviewed, verified, and refined all AI-assisted outputs. All related-work analysis, experimental results, and conclusions were produced independently by the authors, who take full responsibility for the content of this work.

ETHICS STATEMENT

This work studies geometry-processing algorithms for chip-packaging structures. It involves no human subjects, no personally identifiable data, and no new dataset release. Public benchmarks are cited in the text. We note no conflicts of interest or harmful applications beyond ordinary dual-use of metrology software.

REPRODUCIBILITY STATEMENT

All quantitative results are produced by deterministic evaluation scripts on the released point clouds, with one exception. The screened-Poisson rows of Table 5 (Appendix D) come from a multi-threaded third-party solver (not bit-reproducible), so we repeat that reconstruction and report the run-to-run range. SDF training hyperparameters are specified in Table 3 and Section 4.1. Post-processing stages (Section 3.5) are rule-based with thresholds stated in the text. Code, configuration files, and evaluation scripts will be open-sourced upon publication.

REFERENCES

Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In *Proceedings of the Sixth ACM Symposium on Solid Modeling and Applications (SMA)*, pp. 249–266, 2001. doi: 10.1145/376957.376986.

Matan Atzmon and Yaron Lipman. SAL: Sign agnostic learning of shapes from raw data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2565–2574, 2020. doi: 10.1109/CVPR42600.2020.00264.

Yizhak Ben-Shabat, Chamin Hewa Koneputugodage, and Stephen Gould. DiGS: Divergence guided shape implicit neural representation for unoriented point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 19323–19332, 2022. doi: 10.1109/CVPR52688.2022.01872.

Mario Botsch, Leif Kobbelt, Mark Pauly, Pierre Alliez, and Bruno Lévy. *Polygon Mesh Processing*. A K Peters/CRC Press, 2010. doi: 10.1201/b10688.

Alexandre Boulch and Renaud Marlet. POCO: Point convolution for surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6302–6314, 2022. doi: 10.1109/CVPR52688.2022.00620.

Jonathan C. Carr, Richard K. Beatson, Jon B. Cherrie, Tim J. Mitchell, W. Richard Fright, Bruce C. McCallum, and Tim R. Evans. Reconstruction and representation of 3D objects with radial basis functions. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 67–76, 2001. doi: 10.1145/383259.383266.

Chao Chen, Yu-Shen Liu, and Zhizhong Han. GridPull: Towards scalability in learning implicit representations from 3D point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 18322–18334, 2023. doi: 10.1109/ICCV51070.2023.01680.

Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5932–5941, 2019. doi: 10.1109/CVPR.2019.00609.

Zhiqin Chen and Hao Zhang. Neural marching cubes. *ACM Transactions on Graphics*, 40(6):1–15, 2021. doi: 10.1145/3478513.3480518.

Zhiqin Chen, Andrea Tagliasacchi, Thomas Funkhouser, and Hao Zhang. Neural dual contouring. *ACM Transactions on Graphics*, 41(4):1–13, 2022. doi: 10.1145/3528223.3530108.

James R. Clough, Nicholas Byrne, Ilkay Oksuz, Veronika A. Zimmer, Julia A. Schnabel, and Andrew P. King. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12): 8766–8778, 2022. doi: 10.1109/TPAMI.2020.3013679.

David Cohen-Steiner, Pierre Alliez, and Mathieu Desbrun. Variational shape approximation. *ACM Transactions on Graphics*, 23(3):905–914, 2004. doi: 10.1145/1186562.1015817.

Qiujie Dong, Rui Xu, Pengfei Wang, Shuangmin Chen, Shiqing Xin, Xiaohong Jia, Wenping Wang, and Changhe Tu. NeurCADRecon: Neural representation for reconstructing CAD surfaces by enforcing zero gaussian curvature. *ACM Transactions on Graphics*, 43(4):1–17, 2024. doi: 10.1145/3658171.

Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010. doi: 10.1090/mbk/069.

Philipp Erler, Paul Guerrero, Stefan Ohrhallinger, Niloy J. Mitra, and Michael Wimmer. Points2Surf: Learning implicit surfaces from point clouds. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 108–124, 2020. doi: 10.1007/978-3-030-58558-7_7.

Rickard Brüel Gabrielsson, Bradley J. Nelson, Anjan Dwaraknath, and Primoz Skraba. A topology layer for machine learning. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 1553–1563, 2020.

Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 209–216, 1997. doi: 10.1145/258734.258849.

Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11):1309–1331, 2009. doi: 10.1002/nme.2579.

Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pp. 3789–3799, 2020.

Hugues Hoppe. Progressive meshes. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 99–108, 1996. doi: 10.1145/237170.237216.

Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 71–78, 1992. doi: 10.1145/133994.134011.

Xiaoling Hu, Fuxin Li, Dimitris Samaras, and Chao Chen. Topology-preserving deep image segmentation. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 5658–5669, 2019.

Yixin Hu, Qingnan Zhou, Xifeng Gao, Alec Jacobson, Denis Zorin, and Daniele Panozzo. Tetrahedral meshing in the wild. *ACM Transactions on Graphics*, 37(4):1–14, 2018. doi: 10.1145/3197517.3201353.

Yixin Hu, Teseo Schneider, Bolun Wang, Denis Zorin, and Daniele Panozzo. Fast tetrahedral meshing in the wild. *ACM Transactions on Graphics*, 39(4):1–18, 2020. doi: 10.1145/3386569.3392385.

Jiahui Huang, Zan Gojcic, Matan Atzmon, Or Litany, Sanja Fidler, and Francis Williams. Neural kernel surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4369–4379, 2023. doi: 10.1109/CVPR52729.2023.00425.

Alec Jacobson, Ladislav Kavan, and Olga Sorkine-Hornung. Robust inside-outside segmentation using generalized winding numbers. *ACM Transactions on Graphics*, 32(4):1–12, 2013. doi: 10.1145/2461912.2461916.

Tao Ju, Frank Losasso, Scott Schaefer, and Joe Warren. Dual contouring of hermite data. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 339–346, 2002. doi: 10.1145/566570.566586.

Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics*, 32(3):1–13, 2013. doi: 10.1145/2487228.2487237.

Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing (SGP)*, pp. 61–70, 2006. doi: 10.2312/SGP/SGP06/061-070.

Leif P. Kobbelt, Mario Botsch, Ulrich Schwanecke, and Hans-Peter Seidel. Feature sensitive surface extraction from volume data. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 57–66, 2001. doi: 10.1145/383259.383265.

Shengtao Li, Ge Gao, Yudong Liu, Ming Gu, and Yu-Shen Liu. Implicit filtering for learning neural signed distance functions from 3D point clouds. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 234–251, 2024. doi: 10.1007/978-3-031-72658-3_14.

William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 163–169, 1987. doi: 10.1145/37401.37422.

Baorui Ma, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. Neural-Pull: Learning signed distance function from point clouds by learning to pull space onto surface. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pp. 7246–7257, 2021.

Nissim Maruani, Roman Klokov, Maks Ovsjanikov, Pierre Alliez, and Mathieu Desbrun. VoroMesh: Learning watertight surface meshes with voronoi diagrams. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 14565–14574, 2023. doi: 10.1109/ICCV51070.2023.01339.

Nissim Maruani, Maks Ovsjanikov, Pierre Alliez, and Mathieu Desbrun. PoNQ: A neural QEM-based mesh representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3647–3657, 2024. doi: 10.1109/CVPR52733.2024.00350.

Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4460–4470, 2019. doi: 10.1109/CVPR.2019.00459.

Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4):1–15, 2022. doi: 10.1145/3528223.3530127.

Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 165–174, 2019. doi: 10.1109/CVPR.2019.00025.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pp. 8024–8035, 2019.

Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 523–540, 2020. doi: 10.1007/978-3-030-58580-8_31.

Songyou Peng, Chiyu Max Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. Shape as points: A differentiable poisson solver. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 13032–13044, 2021.

Ruwen Schnabel, Roland Wahl, and Reinhard Klein. Efficient RANSAC for point-cloud shape detection. *Computer Graphics Forum*, 26(2):214–226, 2007. doi: 10.1111/j.1467-8659.2007.01016.x.

Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: A hybrid representation for high-resolution 3D shape synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 6087–6101, 2021.

Tianchang Shen, Jacob Munkberg, Jon Hasselgren, Kangxue Yin, Zian Wang, Wenzheng Chen, Zan Gojcic, Sanja Fidler, Nicholas Sharp, and Jun Gao. Flexible isosurface extraction for gradient-based mesh optimization. *ACM Transactions on Graphics*, 42(4):1–16, 2023. doi: 10.1145/3592430.

Hang Si. TetGen, a delaunay-based quality tetrahedral mesh generator. *ACM Transactions on Mathematical Software*, 41(2):1–36, 2015. doi: 10.1145/2629697.

Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 7462–7473, 2020.

Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 7537–7547, 2020.

Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 27171–27183, 2021.

Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomír Měch, and Ulrich Neumann. DISN: Deep implicit surface network for high-quality single-view 3D reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 490–500, 2019.

Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 4805–4815, 2021.

Chen Zhang, Wentao Wang, Ximeng Li, Xinyao Liao, Wanjuan Su, and Wenbing Tao. High-fidelity lightweight mesh reconstruction from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11739–11748, 2025. doi: 10.1109/CVPR52734.2025.01096.

Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3D: A modern library for 3D data processing. *arXiv preprint arXiv:1801.09847*, 2018.

Qingnan Zhou and Alec Jacobson. Thingi10K: A dataset of 10,000 3D-printing models. *arXiv preprint arXiv:1605.04797*, 2016. doi: 10.48550/arXiv.1605.04797.

A Additional implementation details

Training budget and cost. SDF training runs per scene at learning rate 10^{-3}: 70K steps for *PackagedChip* (batch 150K, 4-step accumulation), 80K for *Serdes* (GT-SDF weight 25), 100K for *CNT-TSV* (GT-SDF weight 50), and 120K for the remaining six models. Wall-clock is 5.8 h (*CNT-TSV*), 5.7 h (*780 Chipset*), 10.5 h (*Serdes*), about 36 h (*PackagedChip*), and 5.0–7.9 h for the five remaining 8-GPU fields. The bilateral baseline uses 100K base steps plus 20K fine-tuning at batch 8,000 with σ halved.

Optimizer and data-side schedule. Adam uses a 1K-step warm-up, cosine annealing to 5% of the base rate, gradient clipping, and NaN/Inf guards; large clouds use gradient accumulation (batch 150K $\times$ 4). *Axis-stretch compensation* rescales sampling noise anisotropically once the bounding-box aspect ratio exceeds 2.5, preventing under-sampling along the short axis of an elongated package, and *normal-weighted sampling* up-weights points with nearly vertical normals, counteracting the sparser scan density on vertical walls. Monitoring meshes use MC at 512^3 during training and 1024^3 at convergence.

Bilateral filtering objective. Each query q is projected to $\hat{q} = q - f_\theta(q)\,\nabla f_\theta(q)/\|\nabla f_\theta(q)\|$ and filtered against its K_{nn} nearest surface points $\{p_k\}$ with normals $\{n_k\}$ (Li et al., 2024):

$$\mathcal{L}_{\mathrm{filter}} = \sum_k w_k \big(|(\hat{q} - p_k)\cdot n_k| + |(\hat{q} - p_k)\cdot \hat{n}_{\hat{q}}|\big), \tag{4}$$

with weights $w_k = \theta_r(\|\hat{q} - p_k\|)\,\phi_n(\hat{n}_{\hat{q}}, n_k)$ at spatial bandwidth $\sigma = 2\,\bar{d}_{1\mathrm{nn}}$, twice the mean 1-NN spacing of the input cloud, and normal width 15°; the filter acts on zero and general level sets alike. Chamfer pulling $\mathcal{L}_{\mathrm{rep}} = \lambda_{cd}\,\mathrm{CD}_{L1}(P_{\mathrm{near}}, \hat{Q})$ ramps from 10 to its maximum over steps 5K–10K; $\mathcal{L}_{\mathrm{zero}} = \mathrm{mean}(|f_\theta(P_{\mathrm{near}})|) + 0.5\,\beta\,\mathrm{TopK}(|f_\theta(P_{\mathrm{near}})|)$ adds a top-k component (k=5%, later 10%), protecting small protrusions but eroding them if pushed further.

Sign and far-field terms. Points farther than 5σ from P must have positive SDF ($\mathcal{L}_{\mathrm{sign}} = \mathrm{mean}(\mathrm{ReLU}(-f_\theta))$), with the sampling box padded along thin axes for razor-thin parts. An air-layer term $\mathcal{L}_{\mathrm{air}} = \mathrm{mean}(\mathrm{ReLU}(-f_\theta + 0.005))$ on in-box samples farther than 3σ from the surface and a far-field term $\mathcal{L}_{\mathrm{far}} = \mathrm{mean}(\mathrm{ReLU}(-f_\theta + 0.01))$ on bounding-box faces and corners suppress spurious negative pockets.

Cubical Betti numbers. Topology is read off the raw 512^3 Marching-Cubes occupancy as a cubical complex (Edelsbrunner & Harer, 2010): $\chi = V - E + F - C$ over its vertices, edges, faces and cubes, b_0 counts 26-connected solid components, b_2 counts 6-connected cavities, and $b_1 = b_0 + b_2 - \chi$ counts independent tunnels. Betti numbers are quoted only where the reference passes the resolution gate of Appendix D.

Hash-encoding schedule. The multi-resolution hash table uses 16 levels with 2 features per level and 2^{19} entries; the coarsest level covers the bounding box with 16^3 cells. Hash features are zero-initialized and gated by a global scale warming up linearly over 2K steps after a 10K-step delay, unlocking one finer level every 100 steps. This prevents the high-frequency encoder from disrupting the already-formed low-frequency SDF.

Void/gap sample generation. Free-space samples for equation 1 are generated by GPU ray casting: rays are cast from through-hole axes and inter-layer mid-surfaces toward the visibility hull, and samples with a conservative clearance to the nearest input point are retained together with their signed clearance as d_j. Candidates are first filtered by a winding-number test (Jacobson et al., 2013) so that only points strictly outside the solid survive. Each retained clearance is then shrunk by a fixed retraction margin (15.0 in CNT-TSV units and 0.003 in normalized units; 1.5M candidates, 8-NN clearance estimation, seed 42), so every certified sample lies strictly on the air side even where the nearest-neighbor distance is quantized by the input sampling. The mechanism is built to fail safe. On sparser or noisier inputs the filter admits fewer samples, so protection degrades by covering less of the void, never by certifying the wrong region, matching the sufficiency of Proposition 1 and the coverage limitation stated in Section 5. For CNT-TSV the generator retains 80,146 void and 46,028 gap samples (126,174 in total). On these samples the depth weight w_j^d of equation 1 concentrates the penalty deep inside the hole, at peak sealing pressure. Each model requires this generation only once (about one minute on GPU).

Post-processing parameters. Planar sharpening quantizes normals to 18 directions (6 axis-aligned, 12 diagonal, 35° threshold), and proximity-aware region growing dissolves chamfer strips before the clamped re-projection and in-plane Laplacian smoothing. Selective QEM is aggressive on planar regions, conservative near features, and followed by a re-sharpening pass once the budget search converges. Feature snapping fits z-layer planes, cylindrical via bores, and creases as lines

Table 3: Loss terms, weights, and activation schedule of the topology-aware SDF objective.

Term	Weight	Active
$\mathcal{L}_{\text{filter}}$ (zero + level sets)	1.0	always
$\mathcal{L}_{\text{rep}}$ (Chamfer L_1)	λ_{cd}: $10 \to \max$	ramp 5–10K
$\mathcal{L}_{\text{zero}}$ (+ top-k)	1.0	always
$\mathcal{L}_{\text{eik}}$	0.3 / 0.03	always
$\mathcal{L}_{\text{sign}}$	0.1	always
$\mathcal{L}_{\text{void}}$, $\mathcal{L}_{\text{gap}}$	5.0 / 5.0	always
$\mathcal{L}_{\text{air}}$	2.0	step $\geq$ 1K
$\mathcal{L}_{\text{far}}$	0.5	step $\geq$ 10K
$\mathcal{L}_{\text{tv}}$	0.005	step $\geq$ 10K
$\mathcal{L}_{\text{gc}}$	0.05	step $\geq$ 2K

and arcs, restoring the flatness and circularity rounded off by decimation. Aspect-ratio repair targets 4.8 rather than 5 to keep a safety margin through the loop described next.

Aspect-ratio repair loop. The repair loop alternates (i) collapsing edges opposite to the largest angle of any triangle with $\text{AR} > 4.8$, (ii) flipping edges when the flip reduces the local maximum AR, and (iii) one pass of constrained Laplacian smoothing restricted to the offending vertices, with displacement clamped so that snapped analytic features (planes, cylinders, creases) are not violated. The loop terminates when no triangle exceeds the bound or after 10 iterations; on all reported models it converges in at most 4 iterations.

Measured effect of the void/gap loss. Complementing Table 2, we track the controlled 20K-step pair over training. The certified-free-space sealed fraction of the run without $\mathcal{L}_{\text{void}}$ peaks at 1.66% (voids) and 3.88% (gaps) at step 10K before GT-SDF supervision partially recovers it, whereas the protected run stays below 0.29% (voids) and 0.96% (gaps) throughout. On the raw 512^3 occupancy at step 20K the protected run reconstructs 13 independent tunnels with no sealed cavity, while the unprotected run loses four tunnels and develops two fully enclosed cavities, so any fluidic or electrical path traced through it would be wrong. In the full-length (70K-step) production runs, the protected reconstruction preserves all through-hole bores of the reference assembly and hole radii deviate from the reference by less than one voxel at 1024^3.

B GEOMETRY COMPLEXITY OF *PackagedChip*

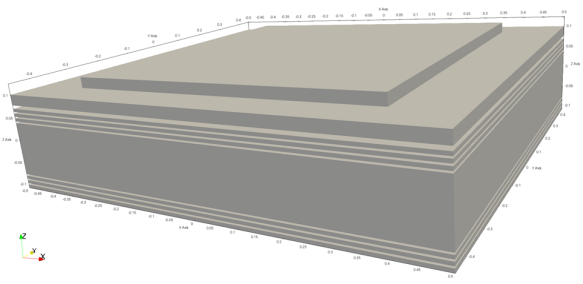

(a) Overall layered structure.

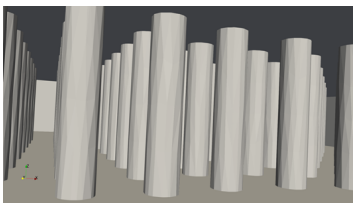

(b) Internal via array.



(c) External layer stack.

Figure 4: **Internal geometry of *PackagedChip*.** (a) The full layered board. (b) Close-up of the via array. (c) Close-up of the thin layer stack. The board packs 591 planes, 96 cylindrical vias, and 103 creases into a slab about one-sixth as thick as its footprint.

PackagedChip of Figure 5 and Table 12 appears as a flat rectangular slab in exterior views. Figure 4 reveals the same 12,674-face mesh from the inside: 591 planes, 96 cylindrical vias, and 103 creases pack into a slab only about one-sixth as thick as its footprint. The vertical edges connecting neighboring layers form via walls; the dense perforations on every layer explain why a pure surface view gives no hint of the geometric complexity inside.

C SURROGATE VERTEX-PLACEMENT BASELINES

The NMC/NDC/VoroMesh entries of Table 1 are re-implementations of the *vertex-placement rule* of each method, evaluated on the frozen SDF field of our trained network so that the placement rule is the only variable. Concretely:

- **NMC surrogate.** The field is sampled on a uniform grid; cells containing a sign change contribute one vertex at the cell centre, with no projection toward the zero level set. This reproduces the defining constraint of grid-bound placement (a vertex cannot leave its cell) without the learned per-configuration vertex offsets of Chen & Zhang (2021)—precisely what the published method contributes.
- **NDC surrogate.** Sign changes are detected on voxel *edges* and each crossing is located by linear interpolation, $t = s_0/(s_0 - s_1)$; the crossings form the vertex set. This is dual contouring's geometric core without the learned vertex prediction of Chen et al. (2022). Because the vertex count equals the number of sign-changing edges, the surrogate cannot be driven to an exact budget and generally returns fewer vertices than requested (Table 1, #V).
- **VoroMesh surrogate.** Seed points are sampled near the surface, their Voronoi vertices are computed, and those with small SDF magnitude are retained as the vertex set. The learned seed optimization of Maruani et al. (2023) is omitted, and a fallback path returns near-surface samples when the Voronoi construction degenerates.

The *MC (equal budget)* and *SPR* entries of Table 1 are not surrogates. MC applies classic Marching Cubes to the frozen field itself and downsamples the result to the requested budget, so its vertices remain locked to voxel edges by construction. SPR fits its own screened-Poisson field to the input cloud through the official Open3D implementation (Zhou et al., 2018), using point normals read from the frozen field's gradient; vertices with frozen-field SDF magnitude exceeding a fixed absolute threshold (0.05) are discarded to remove extrapolated bubble surfaces. The remainder is sampled to the same budget without density trimming, and the octree depth is swept over $\{8, 9, 10\}$ and selected by Chamfer distance against the shared reference, with the full sweep recorded in the backing data. SPR is scored against the same 512^3 reference as every other entry, and it is the only entry with vertices off the frozen field's zero surface. SPR is not bit-reproducible, since the solver is multi-threaded: the table reports one realization per model, and the caption gives the run-to-run spreads measured over three realizations at the selected depth (`results/phase1_extra/*/spr_repeats.json`). The consequences for interpretation are those stated in Section 4.1: these values quantify how much of the fidelity gap is explained by grid-bound placement alone, they are not reproductions of the published systems, and they must not be compared against the numbers reported in the original papers. Nothing here affects the *Ours* rows, produced by the full pipeline, or Table 5, with its competing entry an official implementation run end to end.

D END-TO-END RECONSTRUCTION PROTOCOL

Non-package geometry exercises the placement rule without the corridors and creases the certificates target, so the three Thingi10K figurines of the shared-field protocol are reported here rather than in the main text (Table 4). Those fields were trained without ground-truth normals, a field-side shortfall rather than a placement one; the rows are the frozen pre-retraining artifact `ai_reports/shadow_tab2_thingi.json`, and what changes once the fields are repaired at far-field sign pressure 0.1 and supervision pad 0.10 is stated at the end of this section.

Table 1 holds the field fixed; Table 5 has every method build its own field. Each pipeline receives the same point cloud and is scored against the ground-truth CAD mesh, so the question is no longer how

Table 4: **Vertex placement on non-package geometry.** Three Thingi10K figurines, the same six placement rules and the same 20K requested budget as Table 1, each scored against the 512^3 Marching-Cubes zero surface of one shared frozen field per model (units and best-in-bold convention as there). Ours ranks first on all 18 model–metric cells, by $4.1\times$ on 354371 CD and by $1.8\times$ (F1) and $6.4\times$ ($\mathrm{VSD}_{\max}$) on 44234. Curved surfaces are the mild case for a curvature-adaptive rule, so these cells speak to generalization and carry none of the topological claims.

Dataset	Method	$\mathrm{VSD}_{\mathrm{avg}}\downarrow$	$\mathrm{VSD}_{\max}\downarrow$	CD↓	NC↑	F1↑	ECD↓
Thingi-44234	NMC	91.88	25.93	19.59	0.713	0.172	0.053
	NDC	17.79	9.90	361.47	0.246	0.136	4.128
	VoroMesh	214.86	75.25	77.51	0.665	0.084	0.073
	MC (equal budget)	17.60	9.87	13.09	0.632	0.517	0.111
	SPR	127.75	53.19	37.71	0.742	0.247	0.037
	Ours	**3.01**	**1.54**	**1.52**	**0.785**	**0.919**	**0.005**
Thingi-47984	NMC	92.03	86.52	36.60	0.852	0.111	0.046
	NDC	5.42	4.77	62.66	0.884	0.741	4.684
	VoroMesh	220.56	183.06	106.01	0.865	0.075	0.058
	MC (equal budget)	5.39	8.37	6.26	0.939	0.884	0.188
	SPR	130.02	194.06	91.66	0.765	0.200	0.026
	Ours	**3.99**	**2.57**	**3.04**	**0.971**	**0.902**	**0.011**
Thingi-354371	NMC	97.29	103.14	31.85	0.858	0.154	0.073
	NDC	3.51	6.45	36.82	0.850	0.421	2.171
	VoroMesh	234.58	129.38	120.19	0.633	0.052	0.418
	MC (equal budget)	2.62	9.46	14.77	0.916	0.640	0.835
	SPR	67.55	53.23	892.11	0.650	0.335	7.806
	Ours	**1.72**	**2.60**	**3.61**	**0.961**	**0.735**	**0.026**

well a rule discretizes a given field but whether the reconstructed surface is right. The six package references come from the project's single mesh-preparation pipeline; *Serdes*'s additionally passes a watertight rebuild stage before use. That same solid enters the solver in Appendix I, so its face count matches the *Serdes* reference row above. All references are built independently of any trained field. The two tables therefore have different references and *must not be compared with each other*. Any margin in this table has to be read against two error scales. The first is the sampling floor: resampling a reference against itself already costs $\mathrm{CD} = 0.15$–4.03×10^{-5} across the nine models, and on *Through-Hole*, *Geforce8500*, *780 Chipset*, *PackagedChip* and *Serdes* both pipelines land within 19% of the respective floor (on PackagedChip, for instance, our 4.30×10^{-5} and SPR's 4.51×10^{-5} against a floor of 4.03×10^{-5}), so the CD cells there are marked † and we draw no conclusion from them. The second is that SPR is not reproducible run to run, since Open3D (Zhou et al., 2018) wraps the official multi-threaded Poisson solver, with its octree reduction order following thread scheduling. We reconstructed it three times at the selected depth and measured the spread of every metric, since a claim could rest on any of them (Table 5 caption); the vertex count is stable but ECD moves by 2–138% depending on the model. F1 saturates at the floor (≥ 0.999) on the three board/via models and does not separate the methods there.

With both scales in hand, the six CAD models read consistently. Normal consistency separates ours from SPR everywhere: ours is higher than SPR's on all six, by 7–$50\times$ SPR's own spread on that metric on the original five (0.9896 vs. 0.9854 on Through-Hole, 0.9921 vs. 0.9906 on Geforce8500, 0.9912 vs. 0.9878 on 780 Chipset, 0.921 vs. 0.908 on PackagedChip, 0.9745 vs. 0.9592 on CNT-TSV) and by 0.093 in absolute terms on *Serdes* (0.9578 vs. 0.8651). The curvature-error gap is wider still, at 4.2–13.5 against 21.7–29.7 on the three board/via models and 43.6–46.6 against 114.4–115.2 on *PackagedChip* and *Serdes*, with the floor at zero, since a screened-Poisson field's global smoothing appears as flattened package creases. Edge Chamfer distance is model-dependent: ours wins on the three board/via models (0.004 vs. 0.016, 0.010 vs. 0.025, 0.009 vs. 0.020; each margin 3–$13\times$ SPR's spread on that model) and on *Serdes* (0.020 vs. 0.063), *PackagedChip* is unresolved (1.255 vs. 1.304×10^{-3}, a margin of only $1.8\times$ SPR's spread), and on CNT-TSV SPR beats ours outright (0.152 vs. our 3.18×10^{-3}). A globally smooth surface has few spurious creases to be penalized for, and our raw Marching-Cubes output is not crease-clean before post-processing. On CNT-TSV the remaining separation is coarser and needs no error scale: SPR's CD is $33\times$ larger

Table 5: **End-to-end reconstruction from the same point cloud**, evaluated against the ground-truth CAD mesh. This table is not comparable with Table 1: here the reference is the ground-truth CAD mesh, whereas Table 1 scores against our frozen field's Marching-Cubes zero surface. The *GT vs. GT* row is the sampling noise floor; † marks differences inside that floor or within twice SPR's run-to-run spread. CD in 10^{-5}, ECD in 10^{-3}, CE is the L_1 curvature error, #V and #F in 10^3. ECD is the two-sided mesh variant, distinct from Table 1. Best excluding the noise-floor row in **bold**.

Dataset	Method	CD↓	NC↑	F1↑	ECD↓	CE↓	#V	#F
CNT-TSV	GT vs. GT (floor)	0.336	0.9914	1.000	0.004	5.91	33.6	33.6
	SPR (Kazhdan & Hoppe, 2013)	639.4	0.9592	0.768	0.152	14.71	260.8	521.3
	NDC (Chen et al., 2022)	**0.339**	**0.9907**	**1.000**	**0.006**	**7.63**	12.3	24.2
	Base objective, MC^{512}	20.91	0.9723	0.885	4.712	13.98	327.3	652.9
	Ours, MC^{512}	19.29	0.9748	0.887	7.463	13.30	337.7	673.7
	Ours, production field, MC	19.79	0.9745	0.886	3.178	16.85	1279.3	2555.0
PackagedChip	GT vs. GT (floor)	4.025	0.9731	0.715	0.075	0.00	114.5	230.5
	SPR (Kazhdan & Hoppe, 2013)	4.507	0.9083	0.681	1.304	115.2	4957.4	9922.3
	NDC (Chen et al., 2022)	8.952	0.8293	0.540	1.413	**10.50**	70.1	134.8
	Ours, production field, MC	4.298†	**0.9207**	**0.691**	1.255†	43.61	1829.8	3700.7
Through-Hole	GT vs. GT (floor)	0.153	0.9874	1.000	0.002	0.00	0.3	0.5
	SPR (Kazhdan & Hoppe, 2013)	0.159	0.9854	1.000	0.016	25.21	333.8	667.7
	NDC (Chen et al., 2022)	0.155†	0.9872	1.000	0.002†	**10.67**	6.8	13.7
	Ours, production field, MC	0.154†	**0.9896**	1.000	0.004	13.48	442.0	883.9
Geforce8500	GT vs. GT (floor)	0.687	0.9905	0.999	0.008	0.00	0.2	0.4
	SPR (Kazhdan & Hoppe, 2013)	0.693	0.9906	0.999	0.025	21.66	813.3	1626.8
	NDC (Chen et al., 2022)	0.704†	0.9892	0.999	0.696	**3.35**	18.0	36.0
	Ours, production field, MC	0.687†	**0.9921**	0.999	**0.010**	4.22	1874.5	3749.1
780 Chipset	GT vs. GT (floor)	0.661	0.9894	0.999	0.007	0.00	0.1	0.2
	SPR (Kazhdan & Hoppe, 2013)	0.669	0.9878	0.999	0.020	29.73	660.1	1320.1
	NDC (Chen et al., 2022)	0.682†	0.9871	0.998	0.009†	**1.34**	17.4	35.0
	Ours, production field, MC	0.660†	**0.9912**	0.999	**0.009**	4.68	1803.4	3606.8
Serdes	GT vs. GT (floor)	0.841	0.9589	1.000	0.020	0.00	318.9	639.4
	SPR (Kazhdan & Hoppe, 2013)	0.994	0.8651	0.998	0.063	114.44	1302.6	2624.7
	NDC (Chen et al., 2022)	1.691	0.8386	0.948	0.305	**10.90**	11.3	25.0
	Ours, production field, MC	0.841†	**0.9578**	**1.000**	**0.020**	46.55	44.6	90.6
Thingi-44234	GT vs. GT (floor)	0.443	0.9946	1.000	1.009	0.00	18.3	36.5
	SPR (Kazhdan & Hoppe, 2013)	0.445†	0.9938	1.000	16.312	**8.28**	289.8	579.7
	NDC (Chen et al., 2022)	0.562†	0.9731	0.998	9.201†	10.50	22.4	45.4
	Ours, production field, MC	0.447	**0.9948**	1.000	10.936†	11.77	1770.2	3540.3
Thingi-47984	GT vs. GT (floor)	0.335	0.9966	1.000	0.050	0.00	140.9	281.7
	SPR (Kazhdan & Hoppe, 2013)	0.336†	0.9961	1.000	3.551	**3.26**	319.4	638.8
	NDC (Chen et al., 2022)	0.411†	0.9795	0.999	5.712	15.91	18.2	36.9
	Ours, production field, MC	0.337	**0.9969**	1.000	3.250†	6.02	1307.3	2614.7
Thingi-354371	GT vs. GT (floor)	0.738	0.9939	0.999	5.345	0.00	129.0	257.9
	SPR (Kazhdan & Hoppe, 2013)	0.740†	0.9935	0.999	19.451	**2.17**	176.3	352.7
	NDC (Chen et al., 2022)	0.919†	0.9738	0.994	20.293	12.13	24.9	50.7
	Ours, production field, MC	0.742	**0.9944**	0.999	15.086†	8.90	2819.1	5638.2

than ours and its F1 drops from 0.887 to 0.768, because the deep through-holes and stacked shells fall outside the observed support, into the regime of its smoothness prior.

Element budgets do not line up consistently, with SPR at 0.2–$0.75\times$ our vertex count on the three board/via models and CNT-TSV but $2.7\times$ on PackagedChip (4.96M vs. 1.83M). Three further results run against us, and we state them here. Beyond SPR's CNT-TSV edge cell, its CNT-TSV curvature error is also lower than our production row's (14.7 vs. 16.9; the certificate-controlled 512^3 extraction reaches 13.3). The topological certificate helps CD, NC, F1 and CE relative to the base

objective ($20.91 \rightarrow 19.29 \times 10^{-5}$, $13.98 \rightarrow 13.30$ CE) but worsens ECD ($4.71 \rightarrow 7.46 \times 10^{-3}$): keeping unobserved channels open leaves more thin geometry for the edge metric to disagree about. The certificate controls topology without improving every surface statistic.

The three Thingi models do not reproduce the CAD pattern either, and we report the split as measured. With the normals-guided sign supervision in place our fields converge on all three (zero-shell fraction 0.0000 at the final dump), and both pipelines now sit at the sampling floor on CD (within 1% of the respective floor rows, hence the † marks), with F1 saturated at ≥ 0.999. Normal consistency still goes to ours on all three (0.9948/0.9969/0.9944 against SPR's 0.9938/0.9961/0.9935, margins 12–18$\times$ SPR's run-to-run spread on that metric), while curvature error goes to SPR (2.2–8.3 against our 6.0–11.8): on organic surfaces a globally smoothed field is the right prior, reversing the package-crease case. Our ECD cells read lower on all three (3.3–15.1 against 3.6–19.5×10^{-3}) but the margins, 8–33%, sit inside SPR's 40–138% Thingi spread, so they carry † and support no claim. Re-running the shared-field protocol on these repaired fields leaves our placement rule ahead on 15 of the 18 Thingi cells rather than all 18 (Table 4): the $\mathrm{VSD}_{\max}$ cell of all three models goes to SPR, with its Poisson vertices carrying no outlier tail. SPR is watertight and manifold on all three Thingi models, whereas on the six CAD models it is neither: organic figurines have no unobserved corridors, so its global-closure prior is the right bias there, while on CNT-TSV that same prior is what fails. Finally, self-intersection mixes properties of reference and mesh: the Thingi-44234 and Thingi-47984 ground-truth meshes themselves carry self-intersecting faces, and our own 44234 and 354371 outputs are likewise not self-intersection-free, on meshes 4–$16\times$ denser than SPR's.

The NDC rows answer the remaining protocol question — whether a learned end-to-end pipeline with official pretrained weights beats a classical one on these clouds — and the answer is metric-dependent, so we report it as measured. On Through-Hole, Geforce8500 and 780 Chipset its CD sits within 3.2% of the floor, inside the † regime. On CNT-TSV its surface statistics land on the sampling floor outright (CD 0.339 against 0.336, NC 0.9907 against our 0.9745, F1 1.000 against 0.887, ECD 0.006 against our 3.178×10^{-3}), and its curvature error undercuts both other pipelines on all six CAD models (1.34–10.90 against our 4.22–46.55 and SPR's 14.7–115.2): both are the signatures of a coarse, smoothly interpolated output, bought at a 4–$104\times$ smaller element budget (6.8–70.1K vertices against our production rows). The same 128^3 grid behind that budget sets the failure mode: the corridor throats of CNT-TSV span one to two voxels on our 512^3 reference and hence roughly a quarter to half a cell on its output grid, so its near-floor CNT-TSV numbers are produced by a mesh on a grid too coarse to represent the certified through-holes; surface statistics saturate while topology is unrepresented—the gap instrumented by Section 4.5 and the motivation for this section's resolution gate. Its clearest surface loss is PackagedChip: the 8192-point input cap leaves ≈ 14 points per plate: F1 0.540 against our 0.691, NC 0.829 against 0.921. On the three Thingi models it likewise sits at the CD floor (within 27%) but NC falls below ours on all three (0.973–0.980 against 0.9948–0.9969), curvature error splits (better on 44234, worse on the other two), and 154–184 faces self-intersect per model. Even with the authors' hole-closing post-processing, none of its nine outputs is watertight or manifold. Its margins are deterministic (fixed weights, fixed sub-sampling seed), so only the floor scale applies to them.

This table omits several entries by design, for reasons stated below. Betti numbers are absent from the table. At the 256^3 occupancy resolution affordable for meshes of this size, the two multi-body references collapse to fewer components than they have (44 mesh bodies to 2 voxel components on CNT-TSV, 9 to 1 on PackagedChip), so the rows on those two models fail the resolution gate of our own protocol and their b_1 may not be quoted. The remaining seven references are single bodies with valid gated voxel Betti numbers and are recorded in the backing JSON, but we print no per-row Betti column valid only on the models with topology least at stake; the topological evidence is reported where the resolution is gated instead, in Section 4.5 and Section 4.6. NeurCADRecon is omitted as well, since its official code is not obtainable on this machine, and we substitute no literature-reported value measured under a different protocol. Simulation meshers such as TetWild and fTetWild are likewise omitted: they repair connectivity of an already-given surface and leave its semantics untouched, so they answer a different question than this table. We also show no simulation-ready post-processed mesh: the only multi-domain artifact available for CNT-TSV was produced from the reference tessellation instead of our field, so its near-floor agreement with the reference would be circular, and the two runs starting from our field fail the aspect-ratio gate and, as Section 4.5 reports, the certificate inequality. Finally, the PackagedChip row uses the last standard-resolution field dump of that run; its two later dumps were extracted on a much finer grid (66.4M

faces, $18\times$ the element count of every other row here) and are therefore not comparable within this table. Watertightness follows the same discipline: the CNT-TSV reference is itself not watertight (a property of the reference, not a verdict on any method), and SPR, expected to be watertight by construction, is neither watertight nor manifold on any of the six CAD models at the density threshold keeping its output untrimmed.

E NON-COMPETITION OF THE FILL PENALTY, AND PROOF OF PROPOSITION 1

The proof is elementary by design: the contribution is the premise, per-sample margins and a Lipschitz bound measurable on shipped artifacts, so the guarantee is checkable rather than assumed. Part (i) proves the non-competition property required of a certificate by Appendix K: the fill penalty of equation 1 can only raise the field at certified samples, and its support is disjoint from the data term's. Part (ii) proves Proposition 1 itself and relates the attained margins to the certified clearances.

Part (i). Write $f_j = f_\theta(x_j)$ and $r_j = \mathrm{ReLU}(d_j - f_j)$, so that $\mathcal{L}_{\text{fill}} = \frac{1}{m}\sum_j r_j^2(1 + w_j^d)$. The map $f_j \mapsto r_j^2$ is differentiable everywhere (including at $f_j = d_j$, with both one-sided derivatives vanishing), with

$$\frac{\partial \mathcal{L}_{\text{fill}}}{\partial f_j} = -\frac{2}{m}\, r_j\,(1 + w_j^d) = \begin{cases} -\frac{2}{m}(d_j - f_j)(1 + w_j^d) < 0, & f_j < d_j, \\ 0, & f_j \geq d_j. \end{cases} \tag{5}$$

Hence the gradient of $\mathcal{L}_{\text{fill}}$ with respect to the field value is nonzero exactly on samples exhibiting a sealing tendency ($f_j < d_j$), and there it is strictly negative, so a gradient-descent step moves f_j upward (toward air). In particular the term exerts no force on any sample with $f_j \geq d_j$, so it cannot push the zero level set outward into certified air. The certified samples are separated from the observed surface by the conservative clearance margin $d_j > 0$, so the supports of $\mathcal{L}_{\text{fill}}$ and of the surface-fidelity terms (bilateral filtering, Chamfer pulling, zero-SDF), acting on the observed surface and its immediate neighborhood, are disjoint. The two groups of losses therefore do not compete at the zero level set. The auxiliary L_1 term of equation 1 is two-sided but its subgradient has magnitude at most $\frac{1}{2m}$ per sample and its minimizer $f_j = d_j$ is strictly positive, so it biases the field toward the certified clearance and away from sealing; the one-sided mechanism above is the dominant effect whenever a sealing tendency exists (r_j large).

Part (ii). Let $m_j = f_\theta(x_j) > 0$ be the margin the trained field attains at x_j. If f_θ is L-Lipschitz, then for any y with $\|y - x_j\| < m_j/L$,

$$f_\theta(y) \;\geq\; f_\theta(x_j) - L\|y - x_j\| \;>\; m_j - L \cdot \frac{m_j}{L} \;=\; 0, \tag{6}$$

so $f_\theta > 0$ on the open ball $B(x_j, m_j/L)$. Let $\gamma \subset \bigcup_j B(x_j, m_j/L)$ be a curve traversing a via covered by these balls (guaranteed when consecutive samples along the via axis are spaced closer than $\min_j m_j/L$—enforced by construction in our ray-casting generator). Every point of γ lies in some ball, hence $f_\theta > 0$ along γ and $\gamma \cap \{f_\theta \leq 0\} = \emptyset$. The zero level set extracted by Marching Cubes therefore cannot intersect γ, *i.e.*, the through-hole remains an open channel in the reconstruction. Both quantities in the hypothesis come off the checkpoint: m_j costs one forward evaluation, and L is estimated from the Eikonal residual, so no assumption that the fill penalty reached zero is needed. At the design optimum $\mathcal{L}_{\text{fill}} = 0$ every r_j vanishes, hence $m_j \geq d_j$ and the balls widen to the certified clearances themselves; in practice $\mathcal{L}_{\text{fill}}$ settles at a small positive value instead, and Table 2 reports $m_j > 0$ for 99.7% of void samples.

F PROOF OF THE CERTIFICATE CHAIN UNDER DISCRETIZATION

We use the notation of Proposition 2 (stated in Section 3.5): f_θ is L-Lipschitz, the trained field attains $f_\theta(x_j) \geq m_j > 0$ at the certified samples, $R_j = m_j/L$, the extraction grid has edge length h, and post-processing stage k moves every vertex by at most ε_k.

Step 1 (continuous positivity balls). Proposition 1 gives $f_\theta > 0$ on $B(x_j, R_j)$ with $R_j = m_j/L$, so the continuous zero level set $\{f_\theta = 0\}$ does not meet that ball.

Step 2 (Marching Cubes). Let T be a triangle produced by Marching Cubes and let c be the grid cell containing it; every MC triangle lies in the convex hull of zero crossings of the trilinear interpolant

on the edges of c, hence inside c. A cell producing a triangle contains a sign change, so at least one of its corners, say p, satisfies $f_\theta(p) \le 0$. By Step 1, $p \notin B(x_j, R_j)$, *i.e.* $\|p - x_j\| \ge R_j$. For any $y \in T \subseteq c$ we have $\|y - p\| \le \mathrm{diam}(c) = \sqrt{3}\,h$, and therefore

$$\|y - x_j\| \;\ge\; \|p - x_j\| - \|y - p\| \;\ge\; R_j - \sqrt{3}\,h. \tag{7}$$

The extracted surface thus avoids $B(x_j,\, R_j - \sqrt{3}h)$. We use the full cell diameter $\sqrt{3}h$ in place of the cell circumradius, because the corner witnessing the sign change need not be the one nearest to y; this makes the constant, and hence equation 3, conservative.

Step 3 (post-processing chain). Each stage k maps every vertex v of the current mesh to a point v' with $\|v' - v\| \le \varepsilon_k$, and the stages are applied in sequence, so after all K stages every surviving vertex satisfies $\|v^{(K)} - v^{(0)}\| \le \sum_k \varepsilon_k$ by the triangle inequality. Let $T = \mathrm{conv}(v_1^{(0)}, v_2^{(0)}, v_3^{(0)})$ be an extracted triangle and $T' = \mathrm{conv}(v_1^{(K)}, v_2^{(K)}, v_3^{(K)})$ its image. For a point $y' = \sum_i \lambda_i v_i^{(K)} \in T'$ with $\lambda_i \ge 0$, $\sum_i \lambda_i = 1$, the corresponding point $y = \sum_i \lambda_i v_i^{(0)} \in T$ satisfies

$$\|y' - y\| \;=\; \Big\| \sum_i \lambda_i \big(v_i^{(K)} - v_i^{(0)}\big) \Big\| \;\le\; \max_i \big\|v_i^{(K)} - v_i^{(0)}\big\| \;\le\; \sum_k \varepsilon_k, \tag{8}$$

so the Hausdorff distance between T and T' is at most $\sum_k \varepsilon_k$. Combining with Step 2, every point y' of the output mesh satisfies $\|y' - x_j\| \ge \|y - x_j\| - \|y' - y\| \ge R_j - \sqrt{3}h - \sum_k \varepsilon_k = \rho_j$, positive whenever equation 3 holds. Stages only removing elements (clustering collapses, quadric contractions, edge collapses in the repair loop) are covered by the same bound, since a collapsed vertex is placed within the clamp radius of the vertices being replaced; stages inserting vertices (edge densification, edge flips) place them on the current surface, hence inside triangles already under the bound.

Step 4 (channel preservation). Let γ be a curve traversing a via with $\gamma \subseteq \bigcup_j B(x_j, \rho_j)$, possible when consecutive certified samples along the via axis are spaced closer than $\min_j \rho_j$. By Step 3 no point of γ lies on the output mesh, so γ is a connected subset of the complement of the output surface and its endpoints (on opposite sides of the via, outside the part) lie in the same component of that complement. Hence the channel is open in the output mesh, and this tunnel contributes to b_1 of the output occupancy. The scope of the claim is limited: it certifies the tunnels covered by $\mathcal{C}$ and leaves tunnels or cavities elsewhere unaddressed, so we still measure b_1, b_2 empirically (Section 4.1) instead of inferring them.

Assumptions and failure modes. The argument makes three assumptions. (a) L is not certified: the Eikonal term enforces $\|\nabla f_\theta\| \approx 1$ only in the mean, so we use an empirical bound obtained from the Eikonal residual, and a field with local gradient spikes inside Ω_0 would violate it. (b) The ε_k must be the *implemented* clamps, not nominal tolerances; a stage with displacement limited only by a stopping criterion contributes an unbounded ε_k and invalidates the conclusion. (c) The condition is sufficient, not necessary, and it is governed by $\min_j m_j$: a single certified sample with a small attained margin makes the global inequality fail even when the chain is safe for the remaining samples. We therefore report both the global test and the per-sample satisfaction rate in Section 4.5.

G Measured terms of the certificate chain

Proposition 2 is only useful if every term of equation 3 can be measured on what we actually ship. We therefore instrument three chains jointly covering the two links of the argument, and measure each term separately on the shipped artifacts themselves. Two of the terms come from the trained network: m_j are the field values attained by the trained checkpoint at the 126,174 certified samples of *CNT-TSV* (80,146 void and 46,028 gap samples), and L is bounded by the empirical distribution of $\|\nabla f_\theta\|$ over 10^6 random probes, reported both at the maximum and at the 99th percentile because the Eikonal term only enforces $L \approx 1$ approximately. The other two come from the extracted geometry: h is recovered from the spacing of the extracted mesh itself, and $\sum_k \varepsilon_k$ is measured end-to-end as the one-sided Hausdorff distance from the final mesh (vertices plus 8×10^5 surface samples) to its source zero level set, using exact point-triangle distances. This is tighter than summing the per-stage clamps and needs no instrumentation inside the post-processing code. Chain B has no learned field: its input is the reference CAD mesh, with an exactly 1-Lipschitz distance function, so $L = 1$

Table 6: Measured terms of equation 3 on *CNT-TSV*, all lengths in normalized units. Chain A isolates the Marching-Cubes link ($\sum_k \varepsilon_k = 0$), chain B the post-processing link ($h = 0$, $L = 1$ exactly), chain C is a complete but *exploratory* run retained here as a negative control. "global" is equation 3 with $\min_j R_j$; "per-sample" is the fraction of certified samples satisfying it individually (at $L = L_{p99}$). "intrusion" is the fraction of certified samples found *inside* the output mesh by the generalized winding number (Jacobson et al., 2013), checking the conclusion independently of the hypothesis.

Chain	$\sqrt{3}\,h$	$\sum_k \varepsilon_k$	$\min_j R_j$	**global**	**per-sample**	**intrusion**	b_1 **in** $\rightarrow$ **out**
A: field $\rightarrow$ MC^{512}	3.73×10^{-3}	0	3.23×10^{-6}	no	84.1%	—	$13 \rightarrow$ —
B: CAD mesh $\rightarrow$ post	0	3.98×10^{-4}	3.16×10^{-3}	yes	100%	0/126174	$2 \rightarrow 2$
C: field $\rightarrow$ MC^{1024} $\rightarrow$ post	1.86×10^{-3}	3.53×10^{-2}	3.12×10^{-7}	no	0%	16932/126174	$7 \rightarrow 12$

and $h = 0$ hold without estimation and equation 3 isolates the post-processing link alone. Table 6 collects the result.

Of the three findings below, one is positive and two are negative. First, the post-processing link holds outright: on chain B the accumulated displacement is 3.98×10^{-4} against a certified margin of 3.16×10^{-3}, a slack of $7.9\times$. Equation 3 holds globally and for every one of the 126,174 samples. No sample is intruded, the realized clearance satisfies $\mathrm{dist}(x_j, \mathrm{mesh}) \geq \rho_j$ for all of them, and the Betti numbers of the voxelized occupancy are unchanged end-to-end ($b_0/b_1/b_2 = 2/2/0$ before and after). Second, the global inequality fails on both chains containing a learned field, and it fails for a reason made explicit in equation 3: $\min_j R_j$ is governed by the worst certified sample, and 0.47% of the samples on chain A (598 of 126,174) already have $m_j \leq 0$, *i.e.* the field itself violates the certificate there, driving $\min_j R_j$ to 3.23×10^{-6} regardless of the discretization. Equation 3 is a sufficient condition, so we report it as one: at $L = L_{p99} = 2.04$ it is satisfied individually by 84.1% of the certified samples on chain A (56.0% at the far more pessimistic $L = L_{\max} = 5.05$), and the topological conclusion is verified separately below. Third, the condition does reject bad chains: chain C is a run never shipped, and the inequality flags it by a wide margin ($\sum_k \varepsilon_k$ is two orders of magnitude larger than on chain B). The independent checks agree with the flag: 13.4% of the certified samples end up inside the output mesh and b_1 moves from 7 to 12. A chain violating equation 3 is precisely a chain with output unfit for a solver.

Where the violated samples sit. A per-via breakdown of chain A rules out any single through-path being systematically uncertified. The 13 through-paths are registered geometrically: on the $f_\theta > 0$ occupancy grid at $h/2$ spacing, each path appears as exactly one enclosed air island at one of the two mid-plane slices ($z = -0.0428$ and $z = +0.0638$), 13 islands in total, matching b_1; each void sample is assigned to the nearest island in (x, y). The 598 samples with $m_j \leq 0$ split into 226 inside via catchments and 372 in gap regions; the via share spreads over 5 of 13 catchments, at most 1.7% of any one catchment's samples. Every catchment's median margin stays positive ($+0.0069$ to $+0.0317$) and 8 of 13 catchments are violation-free (Table 7). The certificate is therefore weakest at scattered samples inside otherwise certified corridors, not at any single through-path, consistent with the independently verified $b_1 = 13$ of chain A; hardening the worst catchments means adding certified samples at those throats, not recovering a lost path.

Two independent measurements corroborate the above. The raw Marching-Cubes occupancy of chain A gives $b_1 = 13$, reproducing the through-path count of Table 2 from an independently written cubical-complex implementation. The alignment used before measuring displacement is a pure translation: the fitted extent ratio is 1.0000001 on chain B, so no hidden rescaling absorbs part of $\sum_k \varepsilon_k$. The limitation exposed by this analysis is that chain B certifies the post-processing operator alone, without the composition of our field with it: our stored multi-domain and layered-board post-processing runs were driven from CAD reference meshes, so the end-to-end composition on a production checkpoint remains verified at the per-sample level only (chain A), with the global test still unmet. Closing that gap requires coupling the snapping displacement bound to the certified clearance, left to future work.

Table 7: Per-via breakdown of chain A on *CNT-TSV*. Each through-path is registered as an enclosed air island of the $f_\theta > 0$ grid at its mid-plane slice, and each void sample is assigned to the nearest island in (x, y). $\min_j m_j$ and med m_j are the worst and median attained field values over the catchment's certified samples; "viol." counts samples with $m_j \leq 0$; "chain" is the per-sample satisfaction rate of equation 3 at $L = L_{p99}$. Gap samples (46,028) are not partitioned into catchments; their 372 violated samples are reported in the text.

via	plate z	#samples	$\min_j m_j$	med m_j	viol.	chain
0	-0.0428	10,998	-0.0129	$+0.0174$	58	81%
1	-0.0428	13,566	$+0.0030$	$+0.0317$	0	91%
2	-0.0428	7,488	-0.0036	$+0.0178$	17	77%
3	-0.0428	11,193	$+0.0030$	$+0.0198$	0	84%
4	-0.0428	9,903	$+0.0029$	$+0.0258$	0	89%
5	-0.0428	880	$+0.0029$	$+0.0153$	0	77%
6	-0.0428	253	$+0.0031$	$+0.0069$	0	42%
7	$+0.0638$	3,393	-0.0037	$+0.0079$	57	53%
8	$+0.0638$	1,638	$+0.0026$	$+0.0085$	0	57%
9	$+0.0638$	4,872	-0.0142	$+0.0241$	71	87%
10	$+0.0638$	2,658	-0.0035	$+0.0091$	23	61%
11	$+0.0638$	12,190	$+0.0029$	$+0.0304$	0	93%
12	$+0.0638$	1,114	$+0.0030$	$+0.0118$	0	68%

H Deferred ablation analyses

This appendix collects the controlled ablation tables and their full analyses referenced from Section 4.6.

Table 8 isolates the two key design choices of the meshing stage. **Surface projection is essential**: replacing explicit level-set projection with voxel-based vertex prediction (Variant A) degrades every metric by up to two orders of magnitude (VSD_{avg} $0.60 \rightarrow 36.55$ on Geforce8500), confirming the analysis of Section 1 that voxel-bound vertices cannot track the implicit surface. **Curvature adaptivity trades uniform accuracy for feature accuracy.** Spreading the same budget uniformly, variant B takes the smooth-dominated cells: CD and NC on both models, plus VSD_{avg} and F1 on Geforce8500. The full model wins where creases decide the metric: ECD improves $6.6\times$ on the crease-rich Geforce8500 ($0.164 \rightarrow 0.025$), and $\text{VSD}_{\max}$ improves or ties on both models. On the thin 780 Chipset the ECD cell goes the other way (1.254 vs. 0.052): at a 2.4% bounding-box aspect ratio nearly every face belongs to the layer stack, so concentrating vertices on creases buys little edge coverage there while costing smooth-region resolution; the full model still keeps F1 and VSD_{avg} on that model. This ECD column is the two-sided detected-edge variant on the ablation-era field revision (see the caption); the protocol ECD of Table 1 is the edge-coverage variant on the final fields. For package simulation, with creases and bores carrying the boundary conditions, the full model's trade-off is the desired one; when smooth fidelity is preferred, the curvature weighting can simply be disabled.

Cross-model isolation without GT-SDF supervision. Table 9 repeats the controlled pair on both models with GT-SDF supervision disabled entirely, so that $\mathcal{L}_{\text{void}}/\mathcal{L}_{\text{gap}}$ is the only term acting on certified free space. The two packages fail in different ways. On *CNT-TSV*, with free space dominated by deep vias, the unprotected field converges to a partially sealed state: at 20K, 16.0% of certified gap samples and 4.4% of void samples are classified as solid, the mean clearance error is inflated by 43–96$\times$, and the field penetrates certified air down to $f_\theta = -0.032$ (normalized). With protection, no certified sample is misclassified from step 10K onward, and the field remains strictly positive throughout certified free space. On *PackagedChip*, a package dominated instead by narrow inter-layer gaps, the failure is transient-dominated: without protection the gaps repeatedly seal during optimization (72.8% of certified gap samples satisfy $f_\theta \leq 0$ at step 5K and 71.0% at 15K). Although the gap happens to reopen by the 20K endpoint, 92.9% of gap samples still sit below their certified clearance ($f_\theta < d_j$) versus 46.0% with protection, so the unprotected field retains no margin against resealing. The protected run never exceeds a transient sealed fraction of 10.4% and ends at 0.05%. The wide bores of *PackagedChip* are preserved by both arms (sealed fraction 2.3×10^{-5}),

Table 8: Ablation of the meshing stage (units as in Table 1). Variant A replaces surface projection with voxel-based vertex prediction; Variant B removes curvature-adaptive resampling. $^*\text{VSD}_{\text{avg}}$ of B is marginally lower (0.586 vs. 0.599) on Geforce8500. The Full rows and both variants were measured on the field revision in use when the ablation was run; the Ours rows of Table 1 use the final fields, so cells shared by the two tables differ by field revision, while each A/B comparison above is internal to a single field. The ECD column retains the two-sided detected-edge variant in use when the ablation was run, distinct from the edge-coverage ECD of Table 1.

Data	Variant	VSD_{avg}	VSD_{max}	CD	NC	F1	ECD
Geforce8500	A (voxel)	36.55	16.31	6.99	0.962	0.437	1.542
	B (no curv.)	0.59	1.66	**2.93**	**0.994**	**0.762**	0.164
	Full	**0.60***	**1.48**	3.88	0.986	0.641	**0.025**
780 Chipset	A (voxel)	32.61	48.48	6.35	0.969	0.494	0.509
	B (no curv.)	0.64	**1.40**	**2.84**	**0.995**	0.773	**0.052**
	Full	**0.16**	**1.40**	3.49	0.963	**0.797**	1.254

Table 9: Cross-model isolation ablation of the void/gap protection loss with GT-SDF supervision *disabled entirely*, so that $\mathcal{L}_{\text{void}}/\mathcal{L}_{\text{gap}}$ is the only term acting on certified free space (two controlled 20K-step pairs per model; cf. Table 2, with GT-SDF co-regularizing clearances). Sealed fraction: share of certified free-space samples with $f_\theta \leq 0$ at 20K; peak gap: worst transient gap sealed fraction over the 5/10/15/20K checkpoints; clearance error: $\text{mean}|f_\theta - d|$ (normalized units, $\times 10^{-3}$); CD in 10^{-3}; entries below 0.01% shown as 0.00.

Model	Variant	Sealed frac. ↓ (%)		Peak gap ↓	Clearance err. ↓		CD↓
		void	gap	(%)	void	gap	
CNT-TSV	with $\mathcal{L}_{\text{void}}$	**0.00**	**0.00**	**0.00**	**0.17**	**0.16**	**4.21**
	w/o $\mathcal{L}_{\text{void}}$	4.38	16.03	37.02	7.18	14.98	4.37
PackagedChip	with $\mathcal{L}_{\text{void}}$	0.00	**0.05**	**10.36**	**0.07**	**0.05**	18.64
	w/o $\mathcal{L}_{\text{void}}$	0.00	0.34	72.84	0.23	1.16	**12.46**

whereas on *CNT-TSV* even the void regions degrade without protection, consistent with sealing risk concentrating in narrow, deep free-space corridors with the scarcest surface observations, the regime of the protection term. This regime carries two caveats. First, protection costs some surface accuracy on *PackagedChip* (CD 18.6 vs. 12.5, $\times 10^{-3}$), a price observed neither on *CNT-TSV* (4.21 vs. 4.37) nor in the GT-SDF setting of Table 2. Second, without GT-SDF supervision the raw-MC occupancy is too fragmented for meaningful Betti bookkeeping (155 vs. 16 spurious components on *CNT-TSV* without vs. with protection; $b_2 = 0$ in both arms), so the field-level statistics above constitute the primary evidence here.

Coverage of the ablation study. The ablations above are indexed by stage: the two meshing-stage choices (surface projection, curvature-adaptive resampling) in Table 8; the void/gap certificate with GT-SDF co-supervision in Table 2 and in isolation across two models in Table 9; and the gradient-consistency loss together with its interaction with the encoder decay in Table 14. The top-k zero-SDF term is assessed qualitatively (it protects small protrusions between steps 5–10K, as noted above). We deliberately did not ablate two of the components on their own, and give the reason for each below. The flat-floor backfill is not toggled separately: its benefit is uniform coverage—what the coverage-sensitive cells of Table 1 (F1, VSD) measure against baselines lacking the mechanism. Its safety is inherited from the discretization certificate, since the backfilled points are zero-set points of the frozen field placed under the same displacement clamp (Section 3.5). A toggle would therefore isolate no additional degree of freedom beyond what Table 1 already shows. The 20K vertex budget is likewise not swept: it is the protocol budget requested identically from every method in Table 1, so varying it would change the protocol rather than the method.

Table 10: Feature-preserving lightweighting of a 2048^3 MC extraction (*Geforce8500*). The pipeline removes all non-manifold edges, closes the surface, and improves triangle quality while reducing vertices by $147\times$.

Property	**MC** 2048^3	**Ours (post)**
Vertices	7,551,165	**51,307**
Faces	15,104,254	**102,610**
Non-manifold edges	1,006	**0**
Watertight	no	**yes**
Min. angle (avg)	44.9°	$\mathbf{52.1}^\circ$
Faces $< 10^\circ$	44,270	**7**
Edge-length ratio	1.4×10^9	**10.0**

Table 11: Sensitivity of the sign-supervision hyperparameters on *Thingi-44234*: four 120K-step retrainings spanning the 2×2 neighborhood of the operating point (far-field sign pressure w=0.1, supervision pad 0.10), plus one single-GPU replicate (‡, batch 40000, no data parallelism). All four sweep runs use DDP mode B (world size 8, per-rank batch 5000; effective batch 40000 and learning rate identical to the single-GPU runs); evaluation follows the protocol of Table 5 on the 120K Marching-Cubes mesh (units as there: CD in 10^{-5}, ECD in 10^{-3}, ECD being the two-sided detected-edge variant). b_1 counts independent tunnels on the 256^3 occupancy; the ground-truth genus is 0. *Thingi-44234* is nearly edge-free ($\approx$150 of the 2×10^5 reference samples pass edge detection), so ECD is measured on a tiny subset and is high-variance, the analogue of the † marking in Table 5; we read insensitivity from the remaining columns. The ‡ replicate matches its DDP counterpart within run-to-run noise. Backing data: `tmp_build/sensitivity/sensitivity_scan.json`.

w **/ pad**	**CD**↓	**NC**↑	**F1**↑	**ECD**↓	**CE**↓	b_1
0.10 / 0.10 (operating point)	0.447	0.9948	1.000	10.94	11.77	0
0.05 / 0.05	0.447	0.9948	1.000	11.63	11.90	0
0.05 / 0.20	0.444	0.9950	1.000	1.81	11.64	0
0.20 / 0.05	0.448	0.9949	1.000	10.80	11.84	0
0.20 / 0.20	0.446	0.9951	1.000	10.35	11.65	0
0.05 / 0.05 (single-GPU)‡	0.447	0.9951	1.000	10.21	11.92	0

Sign-supervision hyperparameter sensitivity. We sweep the two operating hyperparameters of the sign-supervision path over the 2×2 neighborhood of the operating point (Table 11). Surface fidelity is flat: CD ranges 0.444–0.448×10^{-5} (spread 0.9%), NC 0.9948–0.9951, F1 sits at ≥ 0.99996, and CE stays within 2.2%; topology is invariant across the sweep (genus 0, watertight, b_1=0). The ECD column is the exception (1.8–11.6×10^{-3}, a $6.4\times$ spread): with only $\approx$150 of the 2×10^5 reference samples passing edge detection on this model, the detected subsets are small and noise-dominated, in line with the 40–138% Thingi ECD spread of the SPR repeats (Table 5), so we do not read a trend into it. Per model, the detected reference edge sets are: under the GT-referenced two-sided protocol of Table 5 (2×10^5 samples), 17,496 / 12,121 / 13,855 / 19,919 / 11,947 / 33,759 on the six packages (Through-Hole, Geforce8500, 780 Chipset, PackagedChip, CNT-TSV, Serdes) and 150 / 503 / 111 on the three Thingi models; under the shared-field protocol of Table 1 (10^5 samples), 4,320 / 4,952 / 5,765 / 13,012 / 5,355 / 6,201 on the same six and 32,658 / 3,938 / 3,127 on the Thingi three. The small-sample caveat is therefore specific to the nearly crease-free Thingi references of the two-sided protocol; the edge-coverage ECD of Table 1 is nowhere in that regime. The single-GPU replicate (batch 40000) of the $(0.05, 0.05)$ corner matches the mode-B DDP run on every column within run-to-run noise (CD 0.447 vs. 0.447, CE 11.92 vs. 11.90), corroborating that the effective-batch-preserving DDP protocol leaves the training outcome unchanged.

Table 12: Simulation-oriented outputs. Left: multi-domain decomposition of *CNT-TSV* (14 labeled domains for multi-material FEM). Right: layered feature-snapped *PackagedChip*. Distances are in normalized units. All six package meshes are carried through to steady-state thermal solves in a multiphysics simulator; the *Geforce8500* lightweight mesh of Table 10 and the watertight *Serdes* package mesh are the two documented in Appendix I.

Property	**CNT-TSV** (multi-domain)	**PackagedChip** (layered)
Output faces	20,056	12,674
Output vertices	10,756	5,577
Domains / primitives	14 domains	591 pl. + 96 cyl.
Creases snapped	–	69 lines + 34 arcs
Chamfer to input	1.26×10^{-5}	–
Hausdorff to input	2.35×10^{-4}	–
Watertight	per-domain	yes
Non-manifold edges	0	0
$\mathrm{AR}_{\mathrm{mean}}$ / AR_{p99}	1.70 / 3.56	1.90 / 4.60
Topology check	genus preserved	pass (tol. 18)
Runtime (s)	119	71

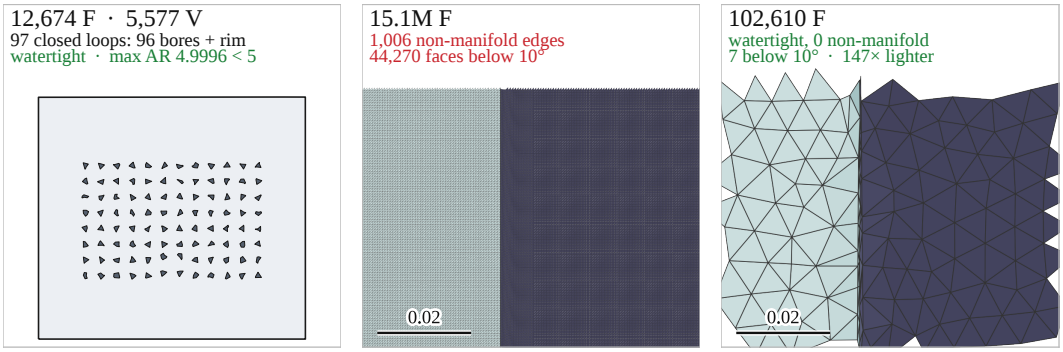

(a) *PackagedChip*, plan cut at $z = -0.007$.

(b) MC 2048^3 before lightweighting.

(c) Ours after lightweighting, same window.

Figure 5: **Output a solver accepts without a manual repair pass.** (a) plan cut of the layered *PackagedChip* at $z = -0.007$: 97 closed loops, 96 bores plus the rim, on a 12,674-face watertight mesh with a worst aspect ratio of 4.9996. (b–c) a window of *Geforce8500* wider than either of Figure 1, before and after lightweighting at one camera: 15.1M faces carrying 1,006 non-manifold edges and 44,270 faces below 10° become 102,610 faces with none and 7, or 248 triangles here instead of 35,872. Remaining columns in Table 12.

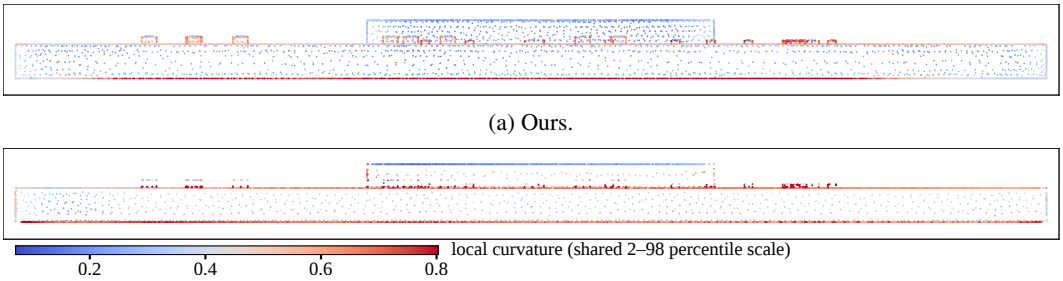

(a) Ours.

(b) Uniform ablation.

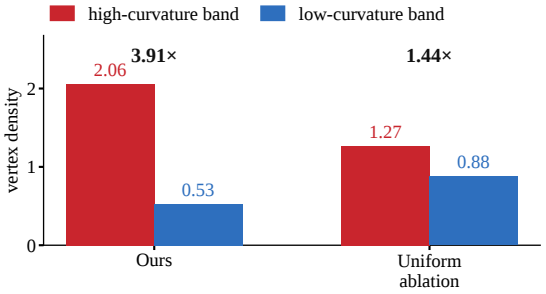

(c) High- vs. low-curvature band densities.

Figure 6: **Curvature-adaptive vertex placement.** Curvature of the vertices placed on *Geforce8500*, (a) ours against (b) the non-adaptive ablation on one shared field and budget, in a cross-section of the board: its z extent is 5.5% of its footprint, so the section is a sliver. Warm points sit on creases, bores, and layer rims; both rows share one 2–98 percentile colour scale. (c) the band densities behind the contrast: $3.91\times$ between the high- and low-curvature bands for ours, against $1.44\times$ without the curvature term.

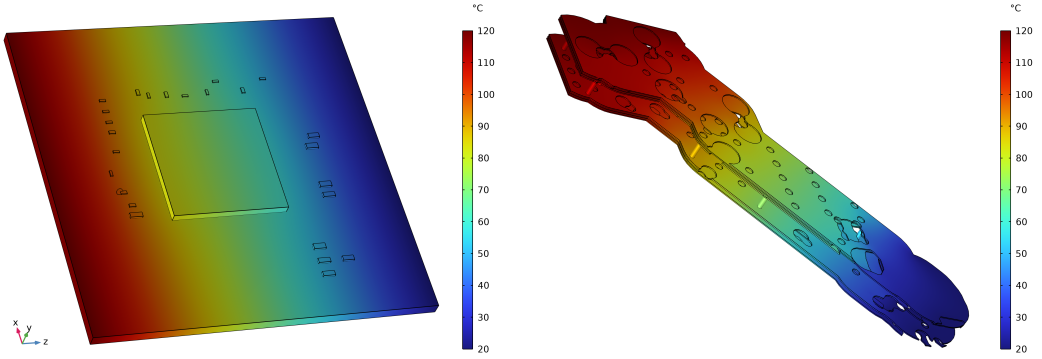

(a) *Geforce8500*. (b) *Serdes*.

Figure 7: Steady-state temperature fields, exported from a multiphysics simulator: (a) the *Geforce8500* pipeline mesh, (b) the *Serdes* pipeline mesh. Copper, fixed boundary temperatures 20 °C and 120 °C, remaining boundaries insulated; both fields span 20–120 °C.

I THERMAL SOLVE ON PACKAGE MESHES

One pipeline produces every package mesh in this paper, and all six package models complete the same path into a multiphysics simulator: automatic import, tetrahedral volume meshing by the automatic mesher, and a stationary heat-conduction solve with no manual geometry editing. Two of them are documented here as representative: for *Geforce8500* the sharpened, simulation-oriented decimation of the production marching-cubes extraction (13,183 triangles), and for *Serdes* the watertight simulation-oriented export of the package (639,412 triangles). The other four (*Through-Hole*, *780 Chipset*, *PackagedChip*, *CNT-TSV*) pass through the same stages and solve the same way. Both STLs arrive as unindexed triangle soups: import stitching merges coincident vertices (39,549 records to 6,594 vertices for *Geforce8500*, 1,918,236 to 318,938 for *Serdes*), after which *Geforce8500* retains 5 open boundary edges that the solver's standard repair closes at its default tolerance, while the *Serdes* export is already watertight with zero non-manifold edges (largest triangle aspect ratio 66.04, circumradius over twice the inradius).

The demonstration physics is stationary heat conduction, the simplest elliptic solve: what a solver demands of its input mesh—watertightness, manifold connectivity, element shape, element budget—is a property of the discretization, not of the PDE, and the Laplace operator used here is the same one as in DC conduction and electrostatics. Coupled multiphysics and transients add solver-side machinery but no new demands on the solid, so one stationary solve per package is enough to certify import-to-solution. Each model is solved as a homogeneous copper block ($k = 400\,\mathrm{W/(m{\cdot}K)}$) with two boundary groups held at 20°C and 120°C and all remaining boundaries thermally insulated. Coordinates are normalized and material constants representative; the demonstration certifies that the mesh carries a solve, not absolute physical quantities.

All six solves complete. The two documented here converge: the tetrahedral discretizations (28,953 and 1,727,111 elements) reach minimum element qualities of 0.136 and 0.010, and the steady-state fields on both packages span 20–120°C, attaining the imposed boundary temperatures—the extrema of a harmonic field sit on the boundary. Figure 7 shows the exported temperature fields; Table 13 lists the import and solve statistics.

Table 13: Mesh import and stationary thermal solve statistics for the two documented solids of Figure 7. Aspect ratio is circumradius over twice the inradius. Both STLs are stored as unindexed triangle soups; after import stitching, *Geforce8500* is meshed after the solver's standard import repair, while the *Serdes* export is already watertight and needs no repair.

Quantity	***Geforce8500***	***Serdes***
surface triangles	13,183	639,412
watertight at meshing	after repair	after stitching
non-manifold edges	0	0
largest aspect ratio ($R/2r$)	5.59	66.04
tetrahedral elements	28,953	1,727,111
minimum element quality (skewness)	0.136	0.010
degrees of freedom	52,358	2,971,797
solve time (s)	1.0	54.0
temperature range (°C)	20–120	20–120

J GRADIENT-CONSISTENCY ABLATION

This appendix states the mechanism behind the multi-scale gradient-consistency loss of equation 2, then collects the full 2×2 ablation table and the residual maps it summarizes (Figure 8).

Trilinear interpolation of hash features is continuous but not continuously differentiable, so the MLP consuming them inherits an axis-aligned lattice of derivative jumps at the cell boundaries of every level; these are what surface as stripes on a planar face. Drawing the offset δ of equation 2 at $1.5\times$ the cell size of a randomly chosen coarse level puts the pair $(x, x+\delta)$ across one of that level's boundaries with probability ≈ 0.98, so the penalty measures exactly those jumps and flattens them. Genuine creases survive: pairs straddling a crease are a measure-zero fraction of the draws, and edge Chamfer distance moves by at most 8% unless both switches are off (Table 14).

Table 14: Controlled 2×2 ablation of the gradient-consistency loss and the per-level encoder decay on the hash-encoded *PackagedChip* model (20K steps per arm). Planar angle: mean normal-angle deviation against the GT plane normal, averaged over the retained large axis-aligned faces (one of the six detected faces is excluded as an evaluation-protocol artifact: it deviates by 77–88° in *all four* arms, independently of either switch). Stripe peak: residual spectral peak at the finest hash level (level 3) along the two in-plane axes (u/v, $\times 10^{-6}$). ECD in 10^{-3} (two-sided detected-edge variant, distinct from the edge-coverage ECD of Table 1). Either mechanism alone suppresses striping; only with both switched off do all three indicators degrade together.

Encoder decay	$\mathcal{L}_{\mathrm{gc}}$	**Planar angle (°)↓**	**Stripe peak↓**	**ECD↓**
on (0.9)	on (0.05)	0.577	4.08/2.10	3.13
on (0.9)	off	0.357	2.37/2.26	3.05
off (1.0)	on (0.05)	0.372	2.78/2.74	2.90
off (1.0)	off	0.744	4.64/4.70	4.61

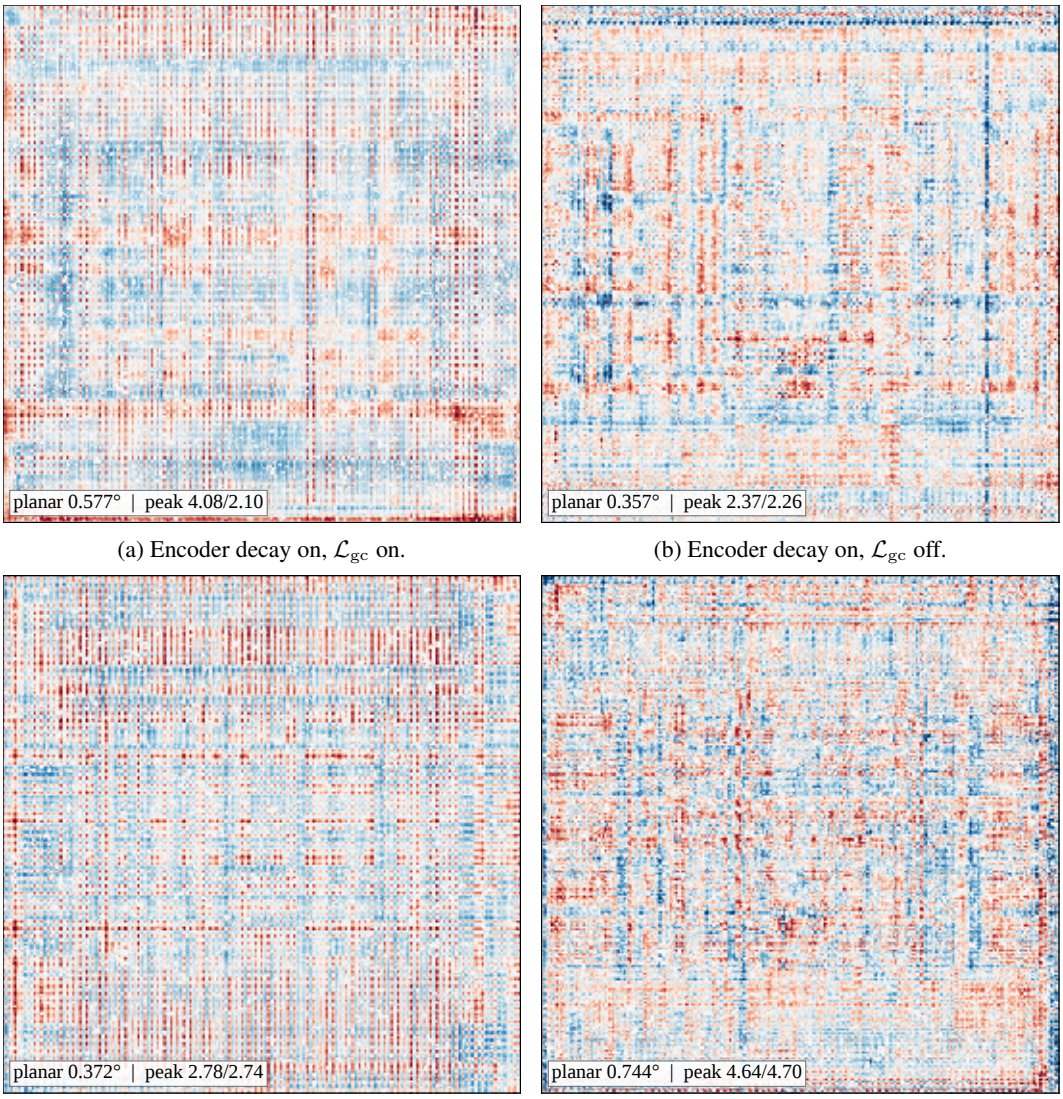

(a) Encoder decay on, $\mathcal{L}_{\mathrm{gc}}$ on.

(b) Encoder decay on, $\mathcal{L}_{\mathrm{gc}}$ off.

(c) Encoder decay off, $\mathcal{L}_{\mathrm{gc}}$ on.

(d) Encoder decay off, $\mathcal{L}_{\mathrm{gc}}$ off.

Figure 8: **The derivative lattice is legible on a planar face, and either mechanism damps it.** Plane residual on the largest axis-aligned face of *PackagedChip* at 20K steps, one panel per arm of Table 14, with that table's planar deviation and level-3 stripe peak printed in each panel. The arm with both switches off (d) carries the densest axis-aligned pattern, the worst deviation ($0.744°$) and the worst peak ($4.64/4.70$). Each panel is scaled to its own $\pm$p99 residual, so the panels compare by pattern rather than by colour magnitude. Composed from the four ablation maps.

K CERTIFICATES INSTEAD OF PRIORS

We constrain the supervision null space Ω_0 of Section 3.2 by objects *verifiable* on the trained field. A *certificate* is a finite family $\mathcal{C} = \{(x_j, d_j)\}_{j=1}^{m}$ with $x_j \in \Omega_0$ and $d_j > 0$, computed *before* training from the input geometry alone; it defines the admissible set $\mathcal{F}(\mathcal{C}) = \{f : f(x_j) \geq d_j \ \forall j\}$ and is enforced by the one-sided penalty

$$\mathcal{L}_{\mathrm{cert}}(f; \mathcal{C}) \;=\; \operatorname*{mean}_{j} \, w_j \, \mathrm{ReLU}\big(d_j - f(x_j)\big)^2, \qquad w_j > 0. \tag{9}$$

A certificate differs from a prior in two properties. *Non-competition*: $\mathrm{supp}(\mathcal{L}_{\mathrm{cert}}) \subseteq \Omega_0$ is disjoint from that of the data term and the penalty is inactive as soon as $f \in \mathcal{F}(\mathcal{C})$, so it trades nothing away on the observed surface (Appendix E, part (i)). *Verifiability*: membership in $\mathcal{F}(\mathcal{C})$ is a finite set of inequalities, checkable on the trained field and, by Proposition 2, on the ultimately simulated mesh.

The three mechanisms of this paper apply the construction of equation 9 to the three failure forms of Section 3.2. The *topological certificate* (equation 1) takes $\mathcal{C}$ to be conservative clearances inside through-holes and inter-layer gaps and yields the per-via guarantee of Proposition 1. The *discretization certificate* (Section 3.5) replaces each post-processing stage's free vertex motion by a hard displacement clamp, transferring the field-level certificate to the output mesh under the single inequality of Proposition 2. The *gradient-consistency regularizer* (equation 2) constrains derivative agreement across the encoder's cell lattice; unlike the other two mechanisms, it is not a certificate in the formal sense (it is enforced in the mean over sampled pairs instead of per sample), and we report it as *measurable but not provable* (Section 4.6 shows it interchangeable with an encoder-side level decay). Existing methods regularize the null space (with smoothness, Eikonal, or curvature priors) but certify nothing in it.

L LIMITATIONS IN DETAIL

(i) Coverage of the certified region. The two certificates and the gradient-consistency regularizer are generated before training, so the method inherits the coverage of the sample generator (Appendix A), and every topology statistic and the shared-field protocol are tied to a 512^3 reference, so structures thinner than a reference voxel are neither certifiable nor measurable.

(ii) Strength of the guarantees. Only the topological guarantee is carried by a proposition, and even Proposition 2 is sufficient rather than necessary, governed by the smallest margin and an approximate Lipschitz constant, so we report per-sample satisfaction (Section 4.5); the gradient-consistency term is enforced in the mean and is interchangeable with an encoder-side decay, so we report it as measured.

(iii) Scope of the end-to-end comparison. The end-to-end comparison (Appendix D) covers all nine models against two external pipelines, one classical and one learned, and its topological column was dropped because multi-body references lose components at the affordable resolution.

(iv) Shape classes gaining least. On organic, curved surfaces the mechanisms buy less, since the topological certificate targets corridors absent from such shapes and feature snapping targets CAD-like primitives; the three figurines nonetheless lead the shared-field protocol, and the end-to-end table favors us once the field is repaired (Appendix D). Optimization is per-scene: every new shape pays the one-off training cost measured in Section 4.1.

(v) Open directions. Three items remain open: priors generalizing across package families, joint certificate-field generation, and clamps coupled to certified clearances so equation 3 holds by construction.