HEATATLAS: SURFACE TEXTURES AS LEARNABLE ATLASES OF LOCAL HEAT KERNELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Heat Kernel Textures (HKTex) represent surface appearance with learnable anisotropic heat kernels, but evaluate each kernel through a global spectral construction that requires fifty Laplace–Beltrami eigendecompositions per mesh and a resident basis of shape $[50, V, 256]$. We show that this evaluator is mismatched to the learned primitive: trained HKTex responses are local and lie in the short-time regime of the heat equation, where the Minakshisundaram–Pleijel expansion is the natural evaluator and depends only on local geometry. HeatAtlas replaces the global eigensystem by radius-bounded local charts from batched mesh unfolding, on which the anisotropic heat kernel becomes a closed-form geodesic Gaussian with a first-order curvature correction. We then make these charts the representation itself: each kernel is a local chart whose non-negative heat response gives a normalised overlap weight and whose appearance is a local function $c_i(u, v)$, instantiated with intrinsic Gabor content in a parallel-transported frame. Separating *where* a primitive contributes from *what* it contains removes the spectral mode ceiling and all global Laplacian preprocessing, and supports learnable per-kernel scale, directional high-frequency appearance, local barriers and edits, planar, disconnected, non-manifold and million-vertex inputs, and transfer across discretisations without retraining. On the 32-object structural-failure set, HeatAtlas reaches 32.29 dB against 19.16 dB for HKTex and 26.82 dB for the strongest baseline, a hash grid; across five meshes of 50k–374k vertices, initialisation is 99–206$\times$ faster and peak memory 20–45$\times$ lower at equal quality; and at matched storage, intrinsic Gabor content improves view PSNR over constant charts by up to $+4.0$ dB on periodic textures, and on real high-frequency assets exceeds the hash grid by $+1.9$ dB (LPIPS $0.044 \rightarrow 0.027$).

1 INTRODUCTION

Surface appearance for 3D content is almost always stored in a UV atlas (Sheffer et al., 2006): a global map to the plane that introduces seams, distortion and uneven texel density, and that binds the texture to one tessellation. Per-element storage such as Ptex and Mesh Colors (Burley & Lacewell, 2008; Yuksel et al., 2010) avoids unwrapping but keeps the binding, while neural textures (Oechsle et al., 2019; Xiang et al., 2021; Koestler et al., 2022; Müller et al., 2022) replace the atlas by a continuous field that is compact but cannot be edited element by element. Heat Kernel Textures (HKTex; Foti et al., 2026) take a third route. A texture is a few thousand anisotropic heat kernels living on the surface, each with a learnable position, orientation, anisotropy, sharpness and colour, optimised intrinsically and densified as in Gaussian splatting (Kerbl et al., 2023). The result is explicit, seam-free and queried directly by a physically based renderer. What limits it is how each kernel is evaluated.

Existing explicit primitives either depend on global state or leave the surface. HKTex evaluates every kernel through truncated eigenfunction expansions of fifty anisotropic Laplace–Beltrami operators (Andreux et al., 2015), a global spectral construction whose basis is computed and kept resident at every vertex. The texture is therefore tied to a mesh-wide state: preprocessing and memory grow with the vertex count, and a fixed number of modes caps the finest detail a kernel can express. Its aligned global construction also assumes a connected manifold with a well-defined principal frame: disconnected components break the shared spectral geometry, exactly planar regions leave the principal direction undefined, and non-manifold soups fall outside the standard cotangent pipeline. Textured and Gabor splats (Rong et al., 2025; Huang & Gong, 2025; Watanabe et al., 2025; Zhou

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.

et al., 2025) remove the global state but leave the surface: their primitives live in Euclidean planes, and Euclidean neighbourhoods couple opposite sides of thin sheets and contacts.

As illustrated in Figure 1, our key observation is that HKTex's thresholded response is effectively local: on a trained spot texture, truncating every response beyond a small radius leaves reconstruction unchanged. Its effective scale lies in the short-time regime of the heat equation, where the Minakshisundaram–Pleijel expansion (Minakshisundaram & Pleijel, 1949; Rosenberg, 1997) is the natural evaluator: a Gaussian in geodesic distance (Varadhan, 1967) with a curvature correction, which depends only on the geometry inside the support. The mode ceiling and the vertex-dependent cost of HKTex both follow from this mismatch. The short-time evaluator needs a local coordinate chart, which we obtain for all kernels at once by unfolding the mesh hinge by hinge into each kernel's tangent plane (Schmidt et al., 2006).

Based on this view, we propose HeatAtlas, which promotes the chart from an evaluation device to the representation itself. Each kernel becomes a learnable local chart. Its non-negative heat response provides a normalised overlap weight that decides *where* the kernel contributes, and its appearance is a local function $c_i(u, v)$ on the chart that decides *what* it contributes. The texture is thus an atlas in the differential-geometric sense, with HKTex as the special case of constant content and spectral weights. We instantiate the content with intrinsic Gabor functions whose frequencies and phases live in the chart's parallel-transported frame, and use learnable patches as a higher-capacity local-content control. The same chart provides per-kernel scale, barriers that unfolding does not cross and local edits, and because the representation depends only on local geometry, it moves between discretisations of a surface without retraining. This separation is an inductive bias for learning signals on surfaces: heat decides where a primitive may act, and the chart supplies the intrinsic coordinates in which its content is learned.

We summarize our contributions as follows: (i) we identify the short-time regime of trained heat-kernel textures and show that a first-order local heat expansion on unfolded charts removes every global eigensystem, lifting the spectral mode ceiling; (ii) we reformulate the texture as a learnable atlas of local functions, in which non-negative heat weights and chart content separate where a primitive contributes from what it contains, and instantiate the content with intrinsic Gabor functions; and (iii) on the HKTex benchmarks, periodic and real high-frequency textures, structural failure cases and scans of up to 2.8M vertices, HeatAtlas improves quality at matched storage, initialises in seconds instead of over an hour, and remains well defined on planar, disconnected, non-manifold and point-sampled inputs.

2 Related Work

2.1 Surface textures and explicit primitives

Textures without a global parameterisation store texels per mesh element (Burley & Lacewell, 2008; Yuksel et al., 2010), represent appearance as a continuous field in 3D or a learned 2D domain (Oechsle et al., 2019; Xiang et al., 2021; Thies et al., 2019; Liang et al., 2026; Koestler et al., 2022; Müller et al., 2022), or diffuse colour over point clouds (Foti et al., 2024). Explicit primitives offer an editable alternative: Gaussian splatting and its surface-aligned variants (Kerbl et al., 2023; Huang et al., 2024; Guédon & Lepetit, 2024; Waczyńska et al., 2027; Gao et al., 2024) carry one colour per primitive, textured splats attach a raster or spatially varying appearance to each one (Chao et al., 2025; Rong et al., 2025; Huang & Gong, 2025; Xu et al., 2024; Younes & Boukhayma, 2026; Papantonakis et al., 2025), and Gabor splats (Wurster et al., 2024; Zhang et al., 2024; Zhu et al., 2025; Watanabe et al., 2025; Zhou et al., 2025; Watanabe & Umetani, 2026) add frequency modulation in the spirit of procedural Gabor noise (Gabor, 1946; Lagae et al., 2009; 2010) and frequency-aware fields (Tancik et al., 2020; Sitzmann et al., 2020; Saragadam et al., 2023; Fathony et al., 2021). HKTex (Foti et al., 2026) is the closest work: its primitives are anisotropic heat kernels on the surface, evaluated through eigen-expansions of anisotropic operators (Andreux et al., 2015). HeatAtlas keeps the explicit, surface-bound primitive of HKTex but gives it a geodesic chart. Unlike splats, whose coordinates live in Euclidean planes and are alpha-composited, our primitives are composed by non-negative normalised weights on the surface itself, and Gabor content is one choice of the chart function $c_i(u, v)$ in a parallel-transported frame rather than a new blending rule.

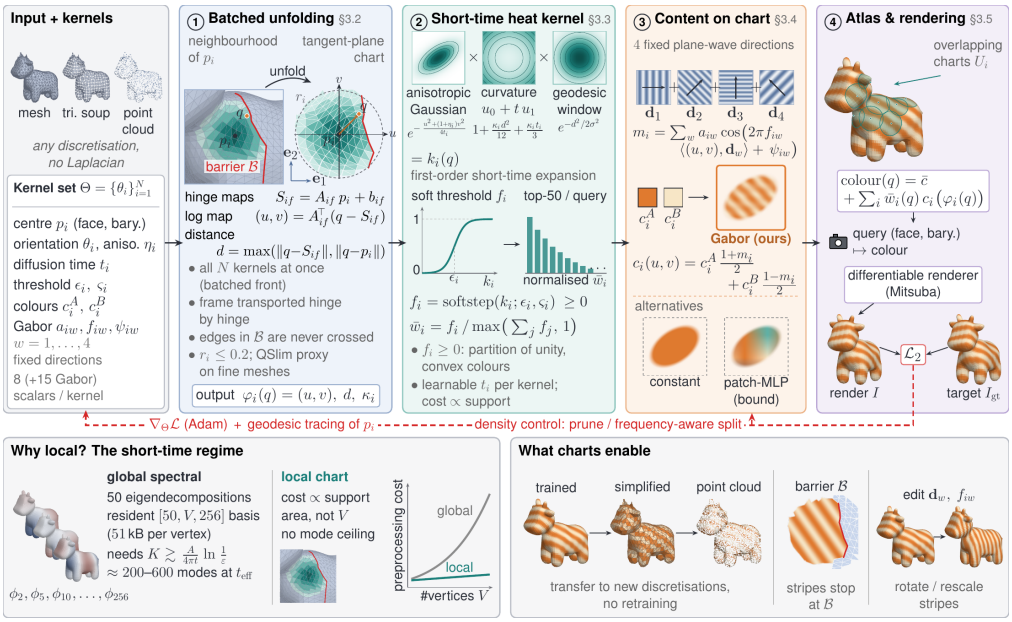


Figure 1: **HeatAtlas overview.** **(1)** The mesh around every centre p_i is unfolded hinge by hinge, giving each query q log-map coordinates (u, v), a distance d and a transported frame; barrier edges $\mathcal{B}$ are never crossed. **(2)** The first-order short-time heat kernel on the chart is soft-thresholded into non-negative overlap weights, normalised to a partition of unity on covered regions. **(3)** Each chart carries content $c_i(u, v)$, by default an intrinsic Gabor pattern. **(4)** Chart contents are composed with unsigned weights, rendered differentiably and optimised end to end. Bottom: the regime argument and what charts enable.

2.2 Heat kernels and local parameterisations

Spectral geometry evaluates heat kernels through eigenpairs of the Laplace–Beltrami operator (Lévy, 2006; Rustamov, 2007; Sun et al., 2009; Sharp et al., 2022), which suits long diffusion times but requires an eigendecomposition per operator and a basis that grows with the mesh. For short times the Minakshisundaram–Pleijel expansion (Minakshisundaram & Pleijel, 1949; Rosenberg, 1997; Varadhan, 1967) describes the kernel locally. Local coordinates around a point can be obtained from discrete or smoothed exponential maps (Schmidt et al., 2006; Herholz & Alexa, 2019), geodesic polar coordinates (Melvær & Reimers, 2012), the vector heat method (Sharp et al., 2019b), intrinsic triangulations (Sharp et al., 2019a; Sharp & Crane, 2020a; Verninas et al., 2026) or exact geodesics (Mitchell et al., 1987; Surazhsky et al., 2005; Crane et al., 2013). We connect these tools through the regime of the learned primitive. Batched hinge unfolding supplies the geometry of a short-time kernel for thousands of moving kernels without assembling a Laplacian, and we make the resulting coordinates the domain of the representation, so the chart determines both where a primitive contributes and what it contains.

3 Method: From local heat kernels to learnable charts

HeatAtlas proceeds in four steps (Figure 1): identify the short-time regime, build one chart per kernel, turn heat responses into non-negative normalised overlap weights, and place a learnable function on each chart.

3.1 Preliminaries: HKTex and the short-time regime

Let $\mathcal{M}$ be a surface given as a triangle mesh, normalised to a bounding-box extent of 2. HKTex represents a texture as N kernels; kernel i has a centre $p_i \in \mathcal{M}$, an in-plane rotation θ_i of the principal-curvature frame, an anisotropy $\eta_i \ge 0$, a soft-threshold level ϵ_i and sharpness ς_i, and a colour c_i. Its response at a query q is an anisotropic heat kernel at a fixed nominal time $t = 1/16$,

evaluated through a truncated eigenfunction series

$$h_t(p,q) \approx \sum_{k=1}^{K} e^{-\lambda_k t}\phi_k(p)\phi_k(q), \qquad K = 256, \tag{1}$$

of the anisotropic Laplace–Beltrami operator (Andreux et al., 2015) stretched by $1+\eta$ along the rotated principal direction. HKTex precomputes aligned eigenpairs for one isotropic and a 7×7 grid of anisotropic operators and interpolates between them. The response is windowed by a Gaussian of width 0.05 in a biharmonic embedding, thresholded, $f_i = \mathrm{softstep}(h_i; \epsilon_i, \varsigma_i) \geq 0$, and the 30 largest of 50 nearest candidates are composed as $\mathrm{colour}(q) = \bar{c} + \sum_i f_i c_i / \max(\sum_i f_i, 1)$. Positions move by geodesic tracing, with error-driven pruning and splitting.

The effective response is local. On a trained HKTex spot texture, zeroing every response beyond radius 0.2 leaves surface PSNR unchanged at 33.62 dB; its fitted effective diffusion time is $t_{\mathrm{eff}} \approx 0.005$, twelve times below the nominal value. HeatAtlas learns this local regime: over $623{,}116$ trained kernels on 130 UV-fitting objects, the radius enclosing 90% of a kernel's contribution is 0.056 at the median and 0.110 at the 99th percentile (Figure 2, middle).

Two evaluators, one regime mismatch. The series equation 1 converges fast for large t; for small t the Minakshisundaram–Pleijel expansion (Minakshisundaram & Pleijel, 1949; Rosenberg, 1997) describes the kernel near the diagonal,

$$h_t(x,y) = \frac{e^{-d(x,y)^2/4t}}{4\pi t}\Big(u_0(x,y) + t\,u_1(x,y) + O(t^2)\Big), \tag{2}$$

with d the geodesic distance (Varadhan, 1967), $u_0 = \Theta_x(y)^{-1/2} = 1 + \frac{\kappa}{12}d^2 + O(d^3)$ for the volume density Θ_x of the exponential map and Gaussian curvature κ, and $u_1(x,x) = \kappa/3$ (Appendix A). The spectral evaluator limits HKTex in two ways. (i) *Mode count.* By Weyl's law $\lambda_K \sim 4\pi K/A$ on a surface of area A, so a relative tail ε in equation 1 at time t needs $K \gtrsim K^* = \frac{A}{4\pi t}\ln\frac{1}{\varepsilon}$; for t_{eff} and $\varepsilon = 10^{-2}$, the benchmark median is $K^* = 437$, with $K^* = 418$ on spot. $K = 256$ lies below the scale of a typical kernel; adding kernels does not help ($+0.5$ dB from 5k to 20k kernels, Table 16), while the local expansion has no mode ceiling. (ii) *Cost.* The series needs 50 eigendecompositions (super-linear in $|\mathcal{V}|$) and a resident $[50, V, 256]$ basis of 51 kB per vertex; on spot preprocessing grows from 122 s at 2.9k to 237 s at 47k vertices. The expansion equation 2 costs only the mesh elements inside a kernel's support and needs d, the direction of y from x, and κ, all of which a local chart supplies.

Atlas view: from support weights to local functions. Differential geometry represents a function on a manifold using charts $\varphi_i : U_i \to \mathbb{R}^2$, a partition of unity ρ_i, and local functions g_i, as $F = \sum_i \rho_i (g_i \circ \varphi_i)$. HeatAtlas makes each part learnable:

$$\mathrm{colour}(q) = \bar{c} + \sum_i \bar{w}_i(q)\; c_i\big(\varphi_i(q)\big), \qquad \bar{w}_i(q) = \frac{f_i(q)}{\max\big(\sum_j f_j(q), 1\big)}. \tag{3}$$

Section 3.2 builds the charts φ_i, Section 3.3 builds the non-negative support weights f_i from equation 2, and Section 3.4 specifies the local functions c_i. Here f_i answers *where* kernel i contributes and $c_i(\varphi_i(q))$ *what* it contributes; HKTex is the special case of constant c_i and spectral f_i.

3.2 Charts by batched radius-bounded unfolding

Radius-bounded unfolding builds the chart around each moving kernel; it supplies both the geometry of the short-time kernel and the domain of local content. Each shared edge is a hinge with a rigid map between the two face planes (Appendix C). All N centres are then developed simultaneously: the state of a reached kernel–face pair (i, f) is a rigid map $S_{if} = A_{if}p_i + b_{if}$, the image of the centre in the plane of f. Pairs advance outward in distance bands, like a batched Dijkstra front, each finalised once (no cycling around cone vertices); a face is retained if its circumcircle comes within the support radius r_i of S_{if} and the segment to its centroid crosses the shared edge (Algorithm 1). Propagating 5,000 kernels takes 0.15–0.6 s on a CPU; supports are rebuilt every 10 steps as centres move.

For a query q on a reached face, the developed displacement rotated back into the source tangent plane and the distance are

$$\delta_{if} = A_{if}^{\mathsf{T}}(q - S_{if}), \qquad d_{if} = \max\big(\|q - S_{if}\|, \|q - p_i\|\big), \tag{4}$$

and the components (u, v) of δ_{if} in the source frame rotated by θ_i define the chart $\varphi_i(q)$. The composition of hinge rotations is a discrete parallel transport, so this frame is consistent across the support up to holonomy. Against exact geodesics on spot the unfolded distance has median relative error 0.03% (p90 1.0%) versus 1.9% (10.8%) for the Euclidean chord, and its direction is within 1.9° of the vector heat log map at the median.

Support, candidates, proxy, barriers. r_i is the distance at which the thresholded response reaches its 1% tail, capped at 0.2. Each query composes at most 50 candidates; on faces supporting more, each query selects its own 50 nearest developed centres, which restores HKTex's per-point semantics on coarse meshes. If the input is much finer than the kernel scale, supports live on a quadric-simplified proxy (Garland & Heckbert, 1997) while queries stay on the input, so support cost depends on kernel footprint, not tessellation. Finally, a set of edges $\mathcal{B}$ that unfolding never crosses (material boundaries, creases, user curves) imposes a no-crossing boundary on the cut surface; normalisation composes kernels entirely from the same side (Appendix A.3), while HKTex would require a new operator and fifty eigendecompositions for each cut.

3.3 Non-negative normalised heat weights

With (u, v), d and a curvature estimate κ_i at the centre (interpolated angle defect), the anisotropic short-time kernel is closed-form. We use its first-order truncation with a geodesic window,

$$k_i(q) = \exp\Big(-\frac{u^2 + (1+\eta_i)v^2}{4t_i}\Big) \underbrace{\Big(1 + \tfrac{\kappa_i}{12} d^2 + \tfrac{\kappa_i}{3} t_i\Big)}_{u_0 + t\, u_1} \exp\Big(-\frac{d^2}{2\sigma^2}\Big), \tag{5}$$

followed by the learned soft threshold $f_i = \mathrm{softstep}(k_i; \epsilon_i, \varsigma_i)$ and the normalisation in equation 3. The first factor is the tangent-plane kernel of the anisotropic operator; the middle factor makes equation 5 a first-order approximation to the surface heat kernel (relative error $O(\kappa d^3) + O(t^2)$); the window $\sigma = 0.1$ reproduces the footprint of HKTex's biharmonic window; with the first factor it sets the effective time $(1/t_i + 2/\sigma^2)^{-1}$ (Section 4.2). We omit the position-independent amplitude $\sqrt{1+\eta_i}/(4\pi t_i)$ because the learned threshold ϵ_i absorbs this per-kernel scale before normalisation. Unlike HKTex, the diffusion time t_i is per-kernel and learnable (log-parameterised, clamped to $[10^{-4}, 1/16]$): under the local evaluator a smaller kernel costs less, and per-kernel scale is what lets quality keep rising with budget on fine textures (Section 4.4).

The heat response serves only as a non-negative support weight. Because $f_i \geq 0$, the weights sum to at most one everywhere and to one wherever $\sum_j f_j \geq 1$. Thus the chart-dependent residual $\sum_i \bar{w}_i c_i(\varphi_i(q))$ is a non-negative combination everywhere and a convex combination of chart contents on covered regions; the final colour adds the global background $\bar{c}$ (Appendix B).

3.4 Content on the chart

Once $\varphi_i(q) = (u, v)$ is available, any compact function on the chart can serve as $c_i(u, v)$. We study three choices. **Constant**, $c_i(u, v) = c_i$, recovers HKTex's one-colour content with our local weights (HeatAtlas (const)) and isolates the effect of the chart. **Intrinsic Gabor** is our main explicit instantiation. Following the finding of Watanabe et al. (2025) that a single learnable wave direction falls into poor minima, it uses $W = 4$ fixed directions $\mathbf{d}_w = (\cos\frac{w\pi}{W}, \sin\frac{w\pi}{W})$ with learnable frequency, phase and amplitude to define an affine modulation of two learned colour vectors:

$$c_i(u, v) = c_i^A \tfrac{1+m_i}{2} + c_i^B \tfrac{1-m_i}{2}, \qquad m_i(u, v) = \sum_{w=1}^{W} a_{iw} \cos\big(2\pi f_{iw} \langle (u, v), \mathbf{d}_w \rangle + \psi_{iw}\big). \tag{6}$$

Frequencies are in cycles per unit geodesic length in the transported frame, independent of tessellation; $c_i^B = c_i^A$ at initialisation so the model starts as HeatAtlas (const). The modulation is chart-local; the non-negativity property applies to the inter-chart weights $\bar{w}_i$.

Oscillation belongs in the content rather than in the heat weight for two reasons. First, in the flat limit a Gaussian times a plane wave is the heat kernel of the magnetic Laplacian $(\nabla - \mathrm{i}\mathbf{k})^2$ (Avron et al., 1978), so a per-kernel frequency is a per-kernel *operator*: unaffordable globally, one cosine on a chart. Second, putting the oscillation in f_i would introduce signed inter-chart coefficients and lose the non-negative hull property of equation 3 (Appendix B).

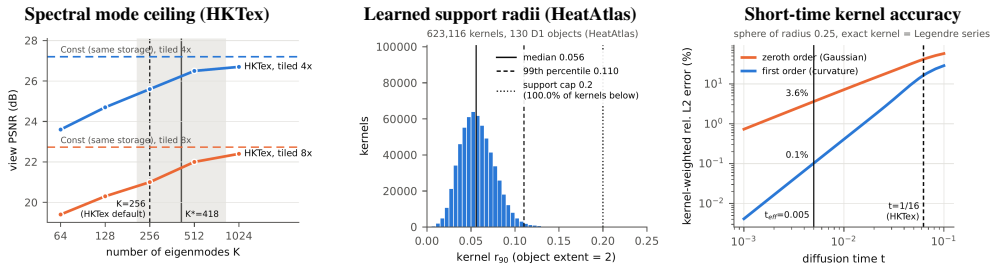


Figure 2: **Locality and the short-time regime.** *Left:* HKTex PSNR vs. eigenmodes K on tiled spot; band: $[0.5, 2] \times K^*$; dashed: HeatAtlas (const) at equal storage. *Middle:* 90% contribution radii of 623,116 trained HeatAtlas kernels on 130 UV-fitting objects. *Right:* error of unwindowed zeroth- and first-order local kernels vs. the exact heat kernel on a sphere of radius 0.25.

Learnable patch (HeatAtlas (patch)) samples a per-kernel $4 \times 4 \times 8$ feature grid at $(u, v)/r_i$ and decodes it with a shared two-layer network. It represents text and decals beyond the Gabor family and serves as a higher-capacity local-content control. A constant kernel stores 8 scalars plus a face index; Gabor adds 15, the patch 128. At matched storage, constant kernels get $2.5\times$ the Gabor count.

3.5 Training, inputs and transfer

For texture fitting the loss is the squared error against a ground-truth texture on area-uniform surface samples with HKTex's importance sampling; for inverse rendering it is an image loss through the Mitsuba differentiable renderer (Jakob et al., 2022). Optimisation, geodesic tracing and density control follow HKTex. For Gabor content a split is made along the dominant wave direction and the children inherit the parent's frequencies (*frequency-aware split*, ablated in Table 5). Unfolding needs only positions, normals, tangent frames and adjacency, which triangle soups, non-manifold meshes and point clouds also provide (Appendix C); no Laplacian is assembled.

A trained texture is a list of centres, local frames and chart parameters. To evaluate it on another discretisation of the *same surface*, we project each centre to the nearest surface point and rebuild the local supports without retraining. Local edits invalidate only the supports within one radius.

4 Experiments

4.1 Setup

Protocol and data. We follow HKTex: five rotating Mitsuba views per object; PSNR, SSIM and LPIPS (Wang et al., 2004; Zhang et al., 2018), with PSNR in view space unless marked as surface PSNR; storage is the zipped size of all parameters. All methods use HKTex's schedule, sampling and density control ($5{,}000$ steps, $1{,}000$ initial kernels, $5\times$ densification cap) on a single GPU (A100 unless noted). We use the 313 UV-fitting and 162 multi-view Objaverse objects of HKTex (Deitke et al., 2023); further sets cover structural failures, 50k–2.8M-vertex meshes, periodic and high-frequency textures, decals, re-discretisations, non-manifold and scanned inputs, and photographs (Appendix D).

Baselines. *HKTex family:* the official code (Foti et al., 2026), HKTex-perpart (one instance per connected component), HKTex-simplify (trained on a $\leq$ 50k-vertex decimation, kernels mapped back by nearest point) and HKTex-K512. *Raster:* the GT texture downsampled to matched storage and vertex colours. *Neural textures:* Intrinsic Neural Fields (Koestler et al., 2022), Texture Fields with Fourier features (Oechsle et al., 2019; Tancik et al., 2020) and a hash grid (Müller et al., 2022). *Primitives on the mesh:* GStex (Rong et al., 2025), 3D Gabor Splatting (Watanabe et al., 2025) and 3DGabSplat (Zhou et al., 2025), constrained to the surface (Appendix E). Our variants are HeatAtlas (const), HeatAtlas (Gabor) (proposed) and HeatAtlas (patch) (higher-capacity control).

4.2 Short-time regime and evaluator choice

Across the 130-object UV-fitting subset, trained HeatAtlas kernels have median effective diffusion time 0.0029 (spot: 0.0028), while their 90% contribution radii have median 0.056 and 99th percentile 0.110. Sweeping K for HKTex on spot (Figure 2, left), PSNR rises through the default $K = 256$,

Table 1: **High-frequency texture reconstruction.** View PSNR (dB), reported as mean ± standard deviation over three seeds, at matched storage (±10%). Synthetic scores average spot and torus; Real contains 20 assets (LPIPS: hash grid 0.044, HeatAtlas (const) 0.031, Gabor 0.027); Text/decals: six decal objects. Splatting baselines are mesh-adapted. HeatAtlas (const) uses 2.5× kernels; the patch is an unranked higher-capacity control. Bold/underline: best/second best.

Method	Stripes $f=8$	Stripes $f=16$	Wood $f=8$	Wood $f=16$	Gabor noise $f=16$	Blobs	Real	Text/ decals
HKTex ($K=256$) (Foti et al., 2026)	30.2 ± 0.12	19.7 ± 0.28	31.4 ± 0.07	21.3 ± 0.26	28.8 ± 0.18	30.5 ± 0.03	25.1 ± 0.25	23.6 ± 0.21
HKTex ($K=512$) (Foti et al., 2026)	31.4 ± 0.12	20.6 ± 0.32	32.6 ± 0.08	22.0 ± 0.33	29.6 ± 0.15	30.8 ± 0.03	25.7 ± 0.26	24.2 ± 0.30
GStex (Rong et al., 2025)	34.1 ± 0.13	26.8 ± 0.47	35.9 ± 0.12	27.2 ± 0.30	32.7 ± 0.24	31.1 ± 0.05	26.4 ± 0.27	29.3 ± 0.28
3D Gabor Splatting (Watanabe et al., 2025)	35.6 ± 0.21	28.4 ± 0.51	37.8 ± 0.11	29.5 ± 0.38	34.3 ± 0.30	30.2 ± 0.06	27.1 ± 0.39	24.9 ± 0.44
3DGabSplat (Zhou et al., 2025)	34.9 ± 0.16	30.1 ± 0.59	38.6 ± 0.11	30.4 ± 0.42	35.2 ± 0.26	31.8 ± 0.05	28.3 ± 0.37	$\mathbf{31.5 \pm 0.37}$
Hash-grid (Müller et al., 2022)	35.3 ± 0.12	27.1 ± 0.25	37.3 ± 0.08	28.7 ± 0.23	35.2 ± 0.12	$\mathbf{31.9 \pm 0.04}$	29.0 ± 0.17	29.7 ± 0.20
HeatAtlas (const)	36.4 ± 0.14	28.9 ± 0.38	38.1 ± 0.09	30.9 ± 0.29	35.5 ± 0.16	30.7 ± 0.05	29.7 ± 0.27	24.2 ± 0.30
HeatAtlas (Gabor, one wave)	$\mathbf{37.3 \pm 0.14}$	30.4 ± 0.55	40.2 ± 0.10	31.6 ± 0.39	35.7 ± 0.28	30.4 ± 0.05	29.8 ± 0.33	24.1 ± 0.39
HeatAtlas (Gabor)	37.1 ± 0.19	$\mathbf{31.5 \pm 0.61}$	$\mathbf{42.1 \pm 0.11}$	$\mathbf{32.5 \pm 0.37}$	$\mathbf{36.8 \pm 0.23}$	31.2 ± 0.05	30.9 ± 0.32	24.2 ± 0.47
HeatAtlas (Gabor, Euclidean)	37.2 ± 0.20	$\mathbf{31.5 \pm 0.51}$	$\mathbf{42.1 \pm 0.13}$	32.4 ± 0.43	$\mathbf{36.8 \pm 0.29}$	31.2 ± 0.05	$\mathbf{31.1 \pm 0.31}$	24.3 ± 0.36
HeatAtlas (patch)	39.5 ± 0.15	33.8 ± 0.41	43.6 ± 0.11	34.7 ± 0.43	38.1 ± 0.20	32.4 ± 0.04	33.0 ± 0.34	28.5 ± 0.33

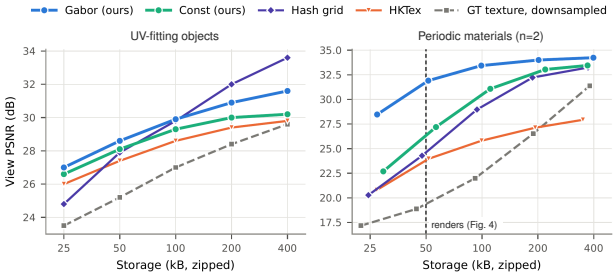


Figure 3: **Equal-storage rate–distortion** on a fixed 30-object UV-fitting subset and two periodic materials.

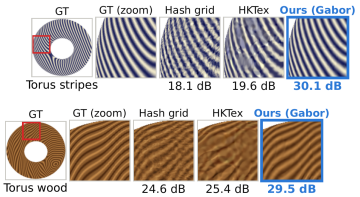


Figure 4: **Matched-storage renders** at ≈ 50 kB on torus stripes and wood.

gains ≈ 1 dB from 256 to 512 modes on the 8× tiling and flattens beyond the predicted $K^* = 418$; even at $K = 1024$ it remains below HeatAtlas (const) at equal storage for both tilings. At the trained scale $t_{\text{eff}} = 0.005$, the zeroth-order Gaussian deviates from the exact kernel by 3.6% and the first-order kernel by 0.1% (Figure 2, right). The curvature correction adds 0.22 dB (Table 9). The learned scale ignores the window: median t_i is 0.0052–0.0064 for $\sigma \in \{0.05, 0.1, 0.2, \infty\}$ (Table 10).

4.3 Quality on the HKTex benchmarks

Replacing the spectral evaluator by charts improves quality on the setting for which HKTex was designed (Table 8, Appendix F). On the 313-object UV-fitting benchmark, HeatAtlas (const) gains +1.00 dB over reproduced HKTex, and storage-matched HeatAtlas (Gabor) gains +2.14 dB; the additional Gabor content contributes +1.14 dB over constant charts. Against the hash grid, the strongest non-heat-kernel baseline, HeatAtlas (Gabor) is on par in PSNR (33.21 vs. 33.11 dB) and lowers LPIPS from 0.059 to 0.043. The difference arises in part on coarse meshes, where HKTex's linear interpolation of eigenvectors inside large triangles makes kernel boundaries polygonal. In multi-view inverse rendering, HeatAtlas (Gabor) reaches 31.63 dB and 0.046 LPIPS on the multi-view set at 0.68× HKTex's training time, and 29.95 dB on real photographs versus HKTex's 26.69 dB (Table 11); Figure 8 shows the recovered printed detail.

4.4 Charts as local function carriers: high-frequency content

A chart encodes structure inside the support instead of spending extra primitives on it (Table 1). Against constant kernels at 2.5× the count, Gabor content changes PSNR by +0.7/ + 2.6 dB on stripes at f=8/16, +4.0/ + 1.6 dB on wood, +1.3 dB on Gabor noise and +0.5 dB on blobs. On the 20 real high-frequency assets the gain is +1.2 dB and LPIPS drops from 0.031 to 0.027. Geodesic and Euclidean phase agree within 0.1 dB on torus and spot, while they differ across creases (+0.15–0.17 dB on slabs with 90° edges, Table 15). A single wave recovers about half of the full Gabor gain on wood and on stripes at f=16, and all of it on stripes at f=8. Text and decals are not locally periodic: Gabor matches constant content at 24.2 dB, and the patch control adds +4.3 dB to 28.5 dB, close to GStex (29.3 dB).

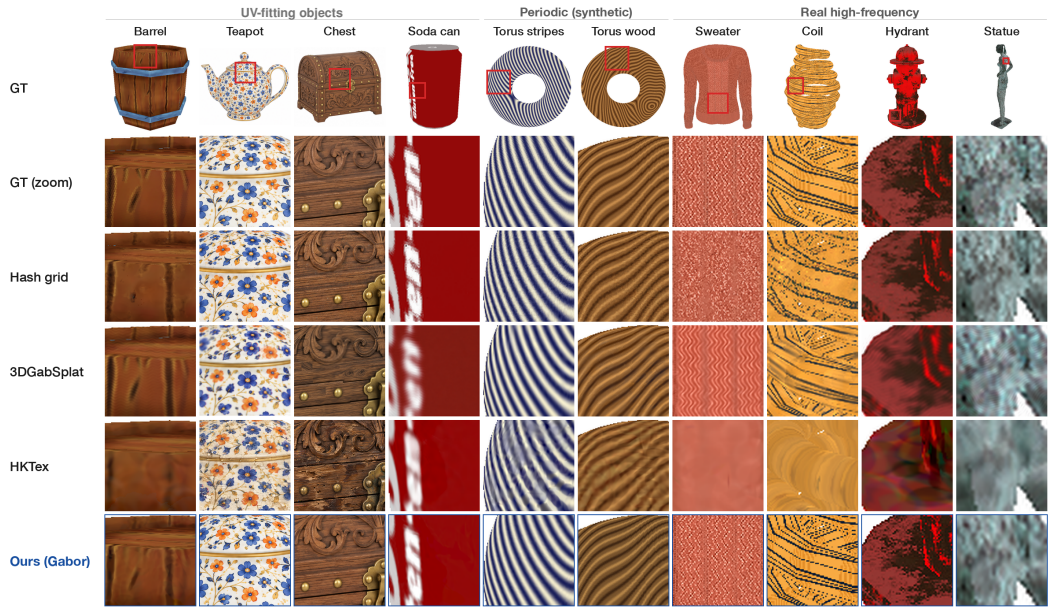


Figure 5: **Reconstruction at equal storage** on UV-fitting objects, synthetic periodic materials and real high-frequency assets: hash grid, 3DGabSplat, HKTex and HeatAtlas (Gabor).

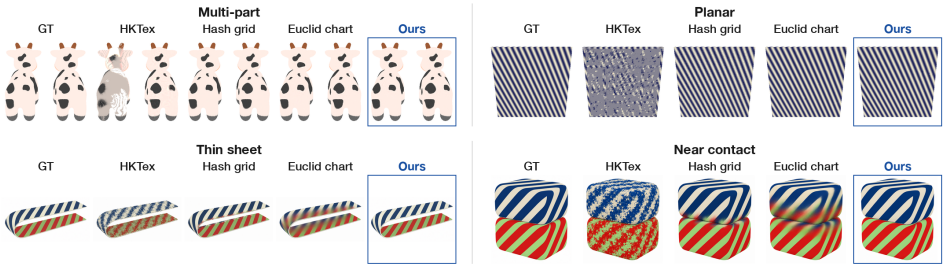


Figure 6: **Structural robustness** on multi-part, planar, thin-sheet and near-contact inputs: HKTex, hash grid, a Euclidean chart (HeatAtlas (Euclid)) and HeatAtlas (const).

Rate–distortion. Chart content helps most at small budgets, where constant kernels cannot afford the primitives an oscillation requires (Figure 3). On the evaluated UV-fitting subset, HeatAtlas (Gabor) exceeds HKTex by 1.0–1.3 dB and the downsampled GT texture by 2.9–3.5 dB between 25 and 100 kB, and the hash grid by 2.2 dB at 25 kB. On the two periodic materials HeatAtlas (Gabor) is the best method at every budget; at 50 kB it reaches ≈ 32 dB against ≈ 24 dB for HKTex and the hash grid, and the torus renders (Figure 4) reach 30.1 vs. 19.6 dB (HKTex) on stripes and 29.5 vs. 25.4 dB on wood. Constant kernels saturate at 250–300 kB, whereas Gabor content still improves at 460 kB.

Qualitatively (Figure 5), HeatAtlas (Gabor) keeps stripe and wood orientation, knit structure and fine carving, where HKTex blurs within each kernel and the hash grid aliases on periodic patterns.

4.5 Robustness, scale and transfer from local charts

Structural failures. Structural failures have two sources (Table 2; per category in Table 7). Multi-part inputs break the shared spectral geometry across components, while exactly planar regions leave the principal frame undefined and collapse the aligned anisotropic basis. Across the complete 32-object structural-failure set, HKTex averages 19.16 dB and HeatAtlas (const) 32.29 dB; the six representative scenes in Table 2 average 19.6 and 33.2 dB. Per-component HKTex recovers only multi-part objects and a robust frame only planar ones (Table 2). Thin sheets and near contacts test intrinsicness: 64% of chord-radius neighbours on a folded ribbon are not intrinsic neighbours. The Euclidean control HeatAtlas (Euclid) (chord distance, tangent projection) ties HeatAtlas (const) on multi-part inputs and is 1.62 dB higher on planar inputs, but loses 11.79 dB on thin sheets and

Table 2: **Six representative structural-failure scenes.** Surface PSNR, reported as mean ± standard deviation over three seeds; the final row is the arithmetic mean of the six scene means. "+frame" is our frame repair of HKTex and ours is HeatAtlas (const). Bold: best per row.

Scene	HKTex	+frame	+perpart +frame	**Ours**
slab, stripes	13.7 ± 0.44	$\mathbf{32.8 \pm 0.28}$	32.6 ± 0.24	30.3 ± 0.28
wavy slab, stripes	16.9 ± 0.61	$\mathbf{32.2 \pm 0.18}$	31.4 ± 0.18	30.5 ± 0.21
two spots	18.4 ± 0.48	18.5 ± 0.47	32.0 ± 0.31	$\mathbf{32.6 \pm 0.30}$
table, 3× subdiv.	40.5 ± 0.13	48.1 ± 0.12	48.1 ± 0.10	$\mathbf{49.1 \pm 0.12}$
multi-part A	16.7 ± 0.61	16.9 ± 0.56	26.2 ± 0.25	$\mathbf{26.7 \pm 0.21}$
multi-part B	11.2 ± 0.65	11.2 ± 0.44	26.7 ± 0.35	$\mathbf{30.0 \pm 0.35}$
mean	19.6 ± 0.24	26.6 ± 0.27	32.8 ± 0.18	$\mathbf{33.2 \pm 0.16}$

Table 3: **Scalability.** Initialisation including preprocessing, peak GPU memory and ΔPSNR (ours − HKTex); mean over all five meshes; full table in Appendix F.

	HKTex		HeatAtlas		
$\lvert\mathcal{V}\rvert$	init	GB	init	GB	ΔPSNR
175k	56 min	18.1	22 s	0.4	+0.5
250k	64 min	25.9	32 s	0.6	0.0
374k	74 min	38.7	45 s	0.9	+0.1
mean	41.8 min	17.8	20.8 s	0.44	+0.11
750k	OOM	–	37 s	0.48	–
0.8–2.8M scans	OOM	–	73–209 s	2.2–8.7	–

4.52 dB at near contacts, where Euclidean neighbourhoods cross the surface (Figure 6). Combining both repairs, HKTex-perpart+frame reaches only 26.62 dB (Table 7).

Scale. Without a resident eigenbasis, geometry cost follows the local support rather than the vertex count (Table 3). On the five 50k–374k meshes, HeatAtlas is within 0.2 dB of HKTex or better while initialising 99–206× faster with 20–45× less memory. Across the 0.8M–2.8M scans, HKTex exhausts an 80 GB A100 at full resolution, whereas HeatAtlas trains on the full meshes with 73–209 s of initialisation and 2.2–8.7 GB, and exceeds HKTex-simplify by 0.57–1.06 dB with sharper relief (Figure 9). With the proxy, step time stays at 17.2 ms at 2.9k and 17.6 ms at 750k vertices.

Transfer, edits and difficult inputs. Re-evaluated on a new discretisation without retraining, a HeatAtlas texture changes view PSNR by −0.63 and −0.82 dB on QEM 25% and 10%, −0.33 dB after isotropic remeshing, +0.00 dB after Loop subdivision and −1.02 dB on the 200k-point cloud. HKTex with a recomputed basis loses 3.75–10.44 dB, while HeatAtlas and the hash grid (−0.03 to −0.49 dB) both transfer as explicit parameter sets (Table 12, Figure 7). Editing a few chart parameters moves, rotates, rescales or re-phases a pattern locally, and a barrier restores a sharp material boundary, raising boundary-band PSNR from 19.08 to 27.25 dB on a coin rim (Figure 10). Table 13 covers non-manifold, soup and scan inputs.

Ablations and efficiency. The largest quality terms are the support choices (Table 5): anisotropy (−7.91 dB on the UV-fitting sample), the geodesic window (−4.24 dB), per-query candidates on coarse meshes and per-kernel scale on fine textures (−2.62 and −2.63 dB). Rebuilding supports every 10 steps instead of every step is 0.26 dB better at 5.6× lower optimisation time. Constant instead of Gabor content at matched storage costs 1.24 dB on UV-fitting and 2.65 dB on stripes. On ten UV-fitting objects, initialisation drops from 163.8 to 4.0 s, peak memory from 2.95 to 0.36 GB, and query throughput rises from 0.15 to 1.44 M points/s (Table 14).

5 Limitations and Conclusion

5.1 Limitations

Density control is error-driven: it refines regions that already carry kernels but cannot seed an isolated region that no kernel reaches, which coverage-driven insertion at small $\sum_j f_j$ would address. Gabor content targets locally periodic, oriented appearance, and text or logos need the learnable patch. The unfolded frame is transported up to holonomy, so Gabor phases around high-curvature cone vertices are ambiguous, and barriers follow mesh edges (Appendix I).

5.2 Conclusion

HeatAtlas turns each one-colour HKTex primitive into a learnable local chart whose closed-form short-time evaluator removes the global spectral pipeline, supplies intrinsic Gabor coordinates and transfers across discretisations without retraining. Evaluated locally, heat-kernel textures scale to million-vertex scans, stay well defined on planar, disconnected and non-manifold inputs, and carry directional high-frequency detail that one-colour primitives cannot afford at matched storage; any compact chart function becomes a texture primitive with unsigned, surface-intrinsic blending.

AI USE STATEMENT

In this work, we used generative AI tools only for editing and polishing the manuscript text. We have not used generative AI tools for generating research ideas, for designing experiments, or for producing or interpreting experimental results, and the remaining disclosure categories are not applicable to this work. All AI-assisted text was reviewed and revised by the authors. We take responsibility for the final content of this work.

ETHICS STATEMENT

This work studies a representation of surface appearance and evaluates it on public 3D asset collections (Objaverse, Thingi10K, Stanford scans, Google Scanned Objects, DTU) under their respective licences; no human subjects, personal data or sensitive attributes are involved. Faster and more robust texture fitting lowers the cost of producing 3D content; we are not aware of specific misuse enabled by this work beyond that of 3D content creation in general. Assets used to construct the failure and high-frequency sets will be listed with identifiers in the code release for licence verification.

REPRODUCIBILITY STATEMENT

The method is fully specified in Section 3 and Appendix C (unfolding algorithm, candidate selection, proxy construction, kernel equation 5, Gabor content equation 6, initialisation, parameter ranges and density control). Appendix A derives the first-order kernel and mode-count bound, and Appendix B proves the non-negative composition property that motivates placing oscillation in the content. Dataset construction is described in Appendix D, baseline adaptations and hyperparameters in Appendix E, and the evaluation protocol in Section 4.1. Upon acceptance, we will release the full source code together with dataset manifests, configurations and scripts for reproducing every table and figure.

REFERENCES

Mathieu Andreux, Emanuele Rodolà, Mathieu Aubry, and Daniel Cremers. Anisotropic Laplace–Beltrami operators for shape analysis. In *Computer Vision – ECCV 2014 Workshops*, pp. 299–312. Springer International Publishing, 2015. doi: 10.1007/978-3-319-16220-1_21.

Joseph Avron, Ira Herbst, and Barry Simon. Schrödinger operators with magnetic fields. I. general interactions. *Duke Mathematical Journal*, 45(4):847–883, 1978. doi: 10.1215/S0012-7094-78-04540-4.

Brent Burley and Dylan Lacewell. Ptex: Per-face texture mapping for production rendering. *Computer Graphics Forum*, 27(4):1155–1164, 2008. doi: 10.1111/j.1467-8659.2008.01253.x.

Brian Chao, Hung-Yu Tseng, Lorenzo Porzi, Chen Gao, Tuotuo Li, Qinbo Li, Ayush Saraf, Jia-Bin Huang, Johannes Kopf, Gordon Wetzstein, and Changil Kim. Textured Gaussians for enhanced 3D scene appearance modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.

Keenan Crane, Clarisse Weischedel, and Max Wardetzky. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics*, 32(5):152:1–152:11, 2013. doi: 10.1145/2516971.2516977.

Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3D objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13142–13153, 2023.

Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B. McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3D scanned household items. In *International Conference on Robotics and Automation*, pp. 2553–2560, 2022.

Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J. Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2021.

Simone Foti, Stefanos Zafeiriou, and Tolga Birdal. UV-free texture generation with denoising and geodesic heat diffusion. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Simone Foti, Caner Korkmaz, Stefanos Zafeiriou, and Tolga Birdal. Heat kernel textures: The geodesic Gaussians that do not splat. In *Computer Vision – ECCV 2026*, pp. 306–323. Springer Nature Switzerland, 2026. doi: 10.1007/978-3-032-37595-7_17. arXiv:2609.07557.

Dennis Gabor. Theory of communication. Part 1: The analysis of information. *Journal of the Institution of Electrical Engineers – Part III: Radio and Communication Engineering*, 93(26): 429–441, 1946. doi: 10.1049/ji-3-2.1946.0074.

Lin Gao, Jie Yang, Bo-Tao Zhang, Jia-Mu Sun, Yu-Jie Yuan, Hongbo Fu, and Yu-Kun Lai. Real-time large-scale deformation of Gaussian splatting. *ACM Transactions on Graphics*, 43(6):1–17, 2024. doi: 10.1145/3687756. arXiv:2402.04796.

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*, pp. 209–216, 1997. doi: 10.1145/258734.258849.

Antoine Guédon and Vincent Lepetit. SuGaR: Surface-aligned Gaussian splatting for efficient 3D mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5354–5363, 2024.

Philipp Herholz and Marc Alexa. Efficient computation of smoothed exponential maps. *Computer Graphics Forum*, 38(6):79–90, 2019. doi: 10.1111/cgf.13607.

Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D Gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, 2024. doi: 10.1145/3641519.3657428.

Zhentao Huang and Minglun Gong. Textured-GS: Gaussian splatting with spatially defined color and opacity. In *Proceedings of the 51st Graphics Interface Conference (GI)*, pp. 1–9, 2025. doi: 10.1145/3769872.3769903. arXiv:2407.09733.

Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. Mitsuba 3 renderer, 2022. `https://mitsuba-renderer.org`.

Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 406–413, 2014.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):139:1–139:14, 2023. doi: 10.1145/3592433.

Lukas Koestler, Daniel Grittner, Michael Möller, Daniel Cremers, and Zorah Lähner. Intrinsic neural fields: Learning functions on manifolds. In *Computer Vision – ECCV 2022*, pp. 622–639. Springer Nature Switzerland, 2022. doi: 10.1007/978-3-031-20086-1_36.

Ares Lagae, Sylvain Lefebvre, George Drettakis, and Philip Dutré. Procedural noise using sparse Gabor convolution. *ACM Transactions on Graphics*, 28(3):54:1–54:10, 2009. doi: 10.1145/1531326.1531360.

Ares Lagae, Peter Vangorp, Toon Lenaerts, and Philip Dutré. Procedural isotropic stochastic textures by example. *Computers & Graphics*, 34(4):312–321, 2010. doi: 10.1016/j.cag.2010.05.004.

Bruno Lévy. Laplace–Beltrami eigenfunctions towards an algorithm that "understands" geometry. In *IEEE International Conference on Shape Modeling and Applications*, pp. 13, 2006. doi: 10.1109/SMI.2006.21.

Yixun Liang, Kunming Luo, Xiao Chen, Rui Chen, Hongyu Yan, Weiyu Li, Jiarui Liu, Fei-Peng Tian, and Ping Tan. UniTEX: Universal high fidelity generative texturing for 3D shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 19917–19927, 2026. arXiv:2505.23253.

Charles Loop. Smooth subdivision surfaces based on triangles. Master's thesis, University of Utah, 1987.

Eivind Lyche Melvær and Martin Reimers. Geodesic polar coordinates on polygonal meshes. *Computer Graphics Forum*, 31(8):2423–2435, 2012. doi: 10.1111/j.1467-8659.2012.03187.x.

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Computer Vision – ECCV 2020*, pp. 405–421, 2020.

Subbaramiah Minakshisundaram and Åke Pleijel. Some properties of the eigenfunctions of the Laplace-operator on Riemannian manifolds. *Canadian Journal of Mathematics*, 1(3):242–256, 1949. doi: 10.4153/CJM-1949-021-5.

Joseph S. B. Mitchell, David M. Mount, and Christos H. Papadimitriou. The discrete geodesic problem. *SIAM Journal on Computing*, 16(4):647–668, 1987. doi: 10.1137/0216045.

Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4):102:1–102:15, 2022. doi: 10.1145/3528223.3530127.

Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3D models, materials, and lighting from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8280–8290, 2022.

Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4531–4540, 2019.

Panagiotis Papantonakis, Georgios Kopanas, Frédo Durand, and George Drettakis. Content-aware texturing for Gaussian splatting. In *Eurographics Symposium on Rendering (Symposium Track)*. The Eurographics Association, 2025. doi: 10.2312/sr.20251190.

Victor Rong, Jingxiang Chen, Sherwin Bahmani, Kiriakos N. Kutulakos, and David B. Lindell. GStex: Per-primitive texturing of 2D Gaussian splatting for decoupled appearance and geometry modeling. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 3508–3518, 2025. arXiv:2409.12954.

Steven Rosenberg. *The Laplacian on a Riemannian Manifold: An Introduction to Analysis on Manifolds*, volume 31 of *London Mathematical Society Student Texts*. Cambridge University Press, 1997.

Raif M. Rustamov. Laplace–Beltrami eigenfunctions for deformation invariant shape representation. In *Eurographics Symposium on Geometry Processing*, pp. 225–233, 2007.

Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G. Baraniuk. WIRE: Wavelet implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18507–18516, 2023.

Ryan Schmidt, Cindy Grimm, and Brian Wyvill. Interactive decal compositing with discrete exponential maps. *ACM Transactions on Graphics*, 25(3):605–613, 2006. doi: 10.1145/1141911.1141930.

Nicholas Sharp and Keenan Crane. You can find geodesic paths in triangle meshes by just flipping edges. *ACM Transactions on Graphics*, 39(6):249:1–249:15, 2020a. doi: 10.1145/3414685.3417839.

Nicholas Sharp and Keenan Crane. A Laplacian for nonmanifold triangle meshes. *Computer Graphics Forum*, 39(5):69–80, 2020b. doi: 10.1111/cgf.14069.

Nicholas Sharp, Yousuf Soliman, and Keenan Crane. Navigating intrinsic triangulations. *ACM Transactions on Graphics*, 38(4):55:1–55:16, 2019a. doi: 10.1145/3306346.3322979.

Nicholas Sharp, Yousuf Soliman, and Keenan Crane. The vector heat method. *ACM Transactions on Graphics*, 38(3):24:1–24:19, 2019b. doi: 10.1145/3243651.

Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. DiffusionNet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics*, 41(3):27:1–27:16, 2022. doi: 10.1145/3507905.

Alla Sheffer, Emil Praun, and Kenneth Rose. Mesh parameterization methods and their applications. *Foundations and Trends in Computer Graphics and Vision*, 2(2):105–171, 2006. doi: 10.1561/0600000011.

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*, volume 33, 2020.

Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, 28(5):1383–1392, 2009. doi: 10.1111/j.1467-8659.2009.01515.x.

Vitaly Surazhsky, Tatiana Surazhsky, Danil Kirsanov, Steven J. Gortler, and Hugues Hoppe. Fast exact and approximate geodesics on meshes. *ACM Transactions on Graphics*, 24(3):553–560, 2005. doi: 10.1145/1073204.1073228.

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*, volume 33, 2020.

Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics*, 38(4):66:1–66:12, 2019. doi: 10.1145/3306346.3323035.

Greg Turk and Marc Levoy. Zippered polygon meshes from range images. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, pp. 311–318, 1994. doi: 10.1145/192161.192241.

S. R. Srinivasa Varadhan. On the behavior of the fundamental solution of the heat equation with variable coefficients. *Communications on Pure and Applied Mathematics*, 20(2):431–455, 1967. doi: 10.1002/cpa.3160200210.

Hippolyte Verninas, Caner Korkmaz, Stefanos Zafeiriou, Tolga Birdal, and Simone Foti. Parallelised differentiable straightest geodesics for 3D meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026.

Joanna Waczyńska, Piotr Borycki, Sławomir Tadeja, Jacek Tabor, and Przemysław Spurek. GaMeS: Mesh-based adapting and modification of Gaussian splatting. *Pattern Recognition*, 181:114561, 2027. doi: 10.1016/j.patcog.2026.114561. arXiv:2402.01459.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.

Haato Watanabe and Nobuyuki Umetani. Neural Gabor splatting: Enhanced Gaussian splatting with neural Gabor for high-frequency surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026. arXiv:2604.15941.

Haato Watanabe, Kenji Tojo, and Nobuyuki Umetani. 3D Gabor splatting: Reconstruction of high-frequency surface texture using Gabor noise. In *Eurographics 2025 – Short Papers*, 2025. arXiv:2504.11003.

Skylar W. Wurster, Ran Zhang, and Changxi Zheng. Gabor splatting for high-quality gigapixel image representations. In *ACM SIGGRAPH 2024 Posters*, 2024.

Fanbo Xiang, Zexiang Xu, Miloš Hašan, Yannick Hold-Geoffroy, Kalyan Sunkavalli, and Hao Su. NeuTex: Neural texture mapping for volumetric neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7119–7128, 2021.

Tian-Xing Xu, Wenbo Hu, Yu-Kun Lai, Ying Shan, and Song-Hai Zhang. Texture-GS: Disentangling the geometry and texture for 3D Gaussian splatting editing. In *Computer Vision – ECCV 2024*. Springer Nature Switzerland, 2024.

Mae Younes and Adnane Boukhayma. TextureSplat: Per-primitive texture mapping for reflective Gaussian splatting. In *International Conference on 3D Vision (3DV)*, 2026. doi: 10.1109/3DV69130.2026.00121. arXiv:2506.13348.

Cem Yuksel, John Keyser, and Donald H. House. Mesh colors. *ACM Transactions on Graphics*, 29 (2):15:1–15:11, 2010. doi: 10.1145/1731047.1731053.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 586–595, 2018.

Xinjie Zhang, Xingtong Ge, Tongda Xu, Dailan He, Yan Wang, Hongwei Qin, Guo Lu, Jing Geng, and Jun Zhang. GaussianImage: 1000 FPS image representation and compression by 2D Gaussian splatting. In *Computer Vision – ECCV 2024*. Springer Nature Switzerland, 2024.

Junyu Zhou, Yuyang Huang, Wenrui Dai, Junni Zou, Ziyang Zheng, Nuowen Kan, Chenglin Li, and Hongkai Xiong. 3DGabSplat: 3D Gabor splatting for frequency-adaptive radiance field rendering. In *Proceedings of the 33rd ACM International Conference on Multimedia (MM)*, pp. 72–81, 2025. doi: 10.1145/3746027.3754855. arXiv:2508.05343.

Qingnan Zhou and Alec Jacobson. Thingi10k: A dataset of 10,000 3D-printing models. *arXiv preprint arXiv:1605.04797*, 2016.

Lingting Zhu, Guying Lin, Jinnan Chen, Xinjie Zhang, Zhenchao Jin, Zhao Wang, and Lequan Yu. Large images are Gaussians: High-quality large image representation with levels of 2D Gaussian splatting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025.

A Derivations

A.1 The short-time expansion and the first-order kernel

Let $(\mathcal{M}, g)$ be a closed Riemannian surface and h_t its heat kernel. For y inside the injectivity radius of x the Minakshisundaram–Pleijel parametrix (Minakshisundaram & Pleijel, 1949; Rosenberg, 1997) gives

$$h_t(x,y) = \frac{1}{4\pi t}\, e^{-d(x,y)^2/4t} \sum_{j=0}^{m} t^j u_j(x,y) + O\big(t^{\,m} e^{-d^2/4t}\big), \tag{7}$$

with $u_0(x,y) = \Theta_x(y)^{-1/2}$, where $\Theta_x(y) = \det(\mathrm{d}\exp_x)$ is the volume density of the exponential map in normal coordinates at x. On a surface with Gaussian curvature κ, $\Theta_x(y) = 1 - \frac{\kappa(x)}{6} d^2 + O(d^3)$, so

$$u_0(x,y) = 1 + \frac{\kappa(x)}{12} d^2 + O(d^3), \qquad u_1(x,x) = \frac{\kappa(x)}{3}, \tag{8}$$

the latter because $u_1(x,x) = R(x)/6$ with scalar curvature $R = 2\kappa$ in two dimensions. On a sphere of radius ρ, $\Theta(d) = \frac{\sin(d/\rho)}{d/\rho}$ exactly, and $u_0 = \big(\frac{d/\rho}{\sin(d/\rho)}\big)^{1/2}$; we use this closed form to validate the implementation (Table 9, Figure 2). The first-order kernel equation 5 keeps u_0 to second order in d and u_1 at the diagonal; its relative error is $O(\kappa d^3) + O(t^2)$ plus the anisotropic correction discussed next. At the representative trained scale $t = 0.005$, the sphere validation measures 0.1% relative error for the first-order kernel.

Anisotropy. HKTex's anisotropic operator is the Laplace–Beltrami operator of the metric $g_D = D^{-1}$ in the rotated principal frame with $D = \mathrm{diag}(1, (1+\eta)^{-1})$, that is, $\Delta_D = \mathrm{div}(D\nabla\cdot)$. In the tangent plane this is the heat kernel with squared distance $u^2 + (1+\eta)v^2$, which is the first factor of equation 5. On a curved surface the metric g_D has its own Gaussian curvature κ_D, which differs from κ by terms involving derivatives of the frame; we use κ in place of κ_D, which is exact when the principal frame is parallel and an $O(\eta\,|\nabla\theta|\,d^2)$ approximation otherwise. On the benchmark, the normalised discrepancy $|\kappa_D - \kappa|/(|\kappa_D| + |\kappa| + \epsilon)$ has median 6.42% and p95 20.15%; replacing κ by κ_D changes view PSNR by only $+0.04$ dB.

Discrete curvature. κ_i at a centre p_i on face f is the barycentric interpolation of the vertex angle defects divided by one third of the incident face areas, computed once per mesh; it is smoothed by one step of uniform Laplacian smoothing to suppress tessellation noise. On the proxy the same quantity is computed on the proxy.

A.2 Mode count bound

By Weyl's law on a closed surface of area A, the counting function satisfies $N(\lambda) \sim \frac{A}{4\pi}\lambda$, so $\lambda_K \sim 4\pi K/A$. The truncation error of equation 1 at the diagonal is $\sum_{k>K} e^{-\lambda_k t}\phi_k(x)^2$; approximating the sum by its Weyl density and the eigenfunction magnitude by its mean $1/A$ gives $\frac{1}{4\pi t} e^{-\lambda_K t}$, that is, a relative tail of $e^{-4\pi K t/A}$ against the leading term $\frac{1}{4\pi t}$. Requiring this to be at most ε yields $K \geq \frac{A}{4\pi t} \ln \frac{1}{\varepsilon}$. The bound is a rule of thumb (it ignores eigenfunction concentration and the aligned interpolation HKTex performs between operators), and Figure 2 (left) is the empirical test.

A.3 Barriers as no-crossing boundaries

Cutting $\mathcal{M}$ along a barrier set $\mathcal{B}$ gives each chart a one-sided domain. Unfolding never crosses the cut, so a kernel contributes only to queries reachable on the same side; the normalisation in equation 3 then composes the available same-side kernels into a partition of unity. This construction enforces sharp material and edit boundaries directly in the chart topology without rebuilding a global operator.

B Non-negative composition

Lemma 1. *Let $f_i \geq 0$ and $\bar{w}_i = f_i / \max(\sum_j f_j, 1)$. Then the chart-dependent residual $r = \sum_i \bar{w}_i\, c_i$ lies in the convex hull of $\{0\} \cup \{c_i\}$, and in the convex hull of $\{c_i\}$ wherever $\sum_j f_j \geq 1$.*

Consequently, the final colour $\bar{c}+r$ lies in the corresponding hull translated by the global background $\bar{c}$.

Proof. $\bar{w}_i \geq 0$ and $\sum_i \bar{w}_i = \min(1, \sum_j f_j) \leq 1$; the deficit $1 - \sum_i \bar{w}_i$ is the weight on 0. □

Remark 1 (Signed weights leave the translated hull)**.** *With two kernels, $f_1 = 1 + \alpha$, $f_2 = -\alpha$ and chart contents $c_1 = 0$, $c_2 = 1$, the normalised residual is $-\alpha/1 = -\alpha$, unbounded as α grows. A Gabor response placed in the weight (a Gaussian times a cosine) changes sign and loses the non-negative hull property. Placing oscillation in the chart content keeps $f_i \geq 0$ and preserves the unsigned inter-chart composition. In a planar toy study (seven synthetic textures, 16–256 kernels), the additive signed-weight variant was 2–4 dB below the non-negative-weight variant on blob and checker textures while the Gabor-content variant was 1–7 dB above constant kernels on the same textures.*

C Implementation details

Unfolding. For each face we precompute its neighbours, centroid and circumradius, normal, and a tangent frame whose first axis follows HKTex's convention (projected mean principal direction, with an edge direction as fallback where the principal frame is undefined); each shared edge is a hinge with a rigid map between the two face planes. Algorithm 1 gives the batched development. Bands are of width equal to the median edge length; a pair (i, f) is finalised when first popped. The edge-window tolerance is 10^{-3} of the edge length. Propagation for 5,000 kernels at radius 0.15 on spot takes 0.575 s single-threaded on CPU and 0.15 s on the proxy; it is rebuilt every 10 steps, which is 0.26 dB better than per-step rebuilding at $5.6\times$ lower optimisation time.

Algorithm 1 Batched radius-bounded unfolding

Input: faces with hinge maps $(R_{f\to f'}, \tau_{f\to f'})$, barrier edges $\mathcal{B}$, centres p_i on faces f_i, radii r_i
Initialise $S_{if_i} \leftarrow p_i$, $A_{if_i} \leftarrow I$; push $(i, f_i, 0)$ into band 0
for band $b = 0, 1, \ldots$ until empty **do**
  for all (i, f) in band b, in parallel **do**
    for all neighbours f' of f with shared edge $e \notin \mathcal{B}$ **do**
      $A_{if'} \leftarrow R_{f\to f'} A_{if}$; $S_{if'} \leftarrow R_{f\to f'} S_{if} + \tau_{f\to f'}$
      $d \leftarrow \|\mathrm{centroid}(f') - S_{if'}\| - \mathrm{circumradius}(f')$
      if $d < r_i$ **and** segment $[S_{if'}, \mathrm{centroid}(f')]$ crosses e **and** (i, f') not finalised **then**
        push (i, f') into band $\lfloor \max(d, 0)/\Delta \rfloor$
      end if
    end for
  end for
end for
Output: finalised pairs $(i, f) \mapsto (A_{if}, S_{if})$

Candidate selection. Faces supporting at most 50 kernels use all of them for every query; on faces supporting more, each query selects its 50 nearest developed centres by d_{if} of equation 4. On the two coarsest benchmark meshes (121 and 330 vertices) a per-face cap left 36.5% and 24.6% of the surface with weight sum below 1; per-query selection reduces this to 1.2% and 1.7% and restores view PSNR from 15.0 to 24.3 and from 30.5 to 48.1 dB (HKTex: 21.5 and 48.1). The cost is a slower step on such meshes (34.6 vs. 20.6 ms).

Proxy. If the median face circumradius is below 0.05 of the median support radius, a QSlim proxy (Garland & Heckbert, 1997) is built with target face size equal to 0.2 of the median support radius; input-face centroids map to proxy faces, supports and candidates live on the proxy, queries and colours on the input; coarse inputs are never refined. On spot subdivided from 2,930 to 749,570 vertices the proxy keeps peak memory at 0.27–0.48 GB and step time at 17.2–17.6 ms.

Parameters and ranges. Per kernel: centre (face index and barycentric coordinates), $\theta_i \in [0, \pi)$, $\eta_i \in [0, 200]$, $\log t_i$ with $t_i \in [10^{-4}, 1/16]$ initialised at $1/16$, threshold ϵ_i and sharpness ς_i in HKTex's ranges, colour $c_i^A \in [0, 1]^3$. Gabor: $c_i^B \in [-1, 1]^3$ initialised to c_i^A plus noise of scale

10^{-2}, $\log f_{iw}$ initialised log-uniformly in $[4, 32]$ cycles per unit length and clamped to $[1, 64]$, $\psi_{iw} \sim \mathcal{U}[0, 2\pi)$, $a_{iw} \sim \mathcal{U}[-0.25, 0.25]$ clamped to $[-1, 1]$. Patch: $G_i \in \mathbb{R}^{4\times4\times8}$ initialised to zero, shared MLP $8 \to 64 \to 64 \to 3$ with ReLU. Geodesic window $\sigma = 0.1$ (selected on spot). A step that produces a non-finite loss or gradient is skipped; this occurred once in $\sim$100 Gabor runs.

Density control. Pruning removes kernels whose accumulated contribution is below HKTex's threshold; error densification splits kernels with the highest accumulated residual, up to $5\times$ the initial count. Frequency-aware split: the child centres are displaced by $\pm\frac{1}{4}$ of the dominant wavelength along $\mathbf{d}_{w^*}$, $w^* = \arg\max_w |a_{iw}|$, and inherit all wave parameters; t_i is halved. For Gabor kernels the contribution statistic uses $|m_i|$-weighted response so that a kernel whose colours nearly cancel is not pruned for being invisible in the mean.

Storage accounting. Storage is the zipped size of a binary file holding, per kernel, the face index (int32), two barycentric coordinates and all scalars in float16, plus a header. For the hash-grid baseline the table and decoder weights are stored in float16 and zipped identically. GT-texture and vertex-colour baselines are stored as PNG and float16 arrays respectively.

Inputs beyond clean meshes. Triangle soups: edges are matched if both endpoints agree within 10^{-5} of the bounding-box extent. Non-manifold edges: each incident face pair is a hinge; the band front finalises whichever development reaches a face first. Point clouds: $k = 16$ nearest neighbours, PCA tangent planes with normals oriented by propagation, hinge between neighbours j, k given by the minimal rotation taking n_j to n_k about $n_j \times n_k$; a "face" is a point with its tangent disc of radius equal to half the mean neighbour distance. Queries on a point cloud are at the points themselves.

D DATASETS

Table 4 summarises the evaluation sets.

Table 4: Evaluation sets used in the experiments.

Content	#	Used in
HKTex UV-fitting set (Objaverse)	313	Tab. 8, 5, Fig. 3
HKTex multi-view set	162	Tab. 11
Structural failures: multi-part, planar, thin sheet, near contact (6 real + 2 synthetic each)	32	Tab. 2, 7
Large meshes, 50k–2.8M vertices	9	Tab. 3, 6
Periodic materials × tilings on spot/torus/slab + real high-frequency assets	48 + 20	Tab. 1, Fig. 3
Text and decals	6	Tab. 1
Re-discretisations of 20 UV-fitting objects	20 × 5	Tab. 12
Non-manifold, triangle soups, scans	5+5+5	Tab. 13
Photographs with meshes	8	Tab. 11

- **UV-fitting and multi-view sets.** The 313 and 162 Objaverse objects of HKTex, selected by its twelve filtering rules (including downloadable and textured, fewer than 60k vertices, static, manifold, one connected component, bounded ray self-intersection count, successful $K = 256$ eigendecomposition, and exclusion of tuning objects). The 130-, 30-, 20- and 10-object subsets are drawn uniformly at random from the 313 UV-fitting objects with a fixed seed; spot, used to select σ, is not among them.
- **Structural failures.** Objects rejected by the "one connected component" rule (multi-part, ≥ 3 components), objects with large exactly planar regions ($\geq 30\%$ of area with zero angle defect), thin double-layer objects (cloth, leaves, packaging with sheet thickness below 0.05), and near-contact objects (two surfaces within 0.05 over $\geq 5\%$ of the area); six real objects per category, plus two synthetic per category: a closed slab and a wavy-top slab with stripes; two spot copies and a slab pair; a thin cylinder shell and a folded ribbon; two tangent spheres and a stacked-plate pair. Synthetic textures are procedural stripes, wood, band-limited noise and blobs at $f \in \{8, 16\}$.
- **Large meshes.** Nine inputs drawn from Thingi10K (Zhou & Jacobson, 2016), the Stanford scanning repository (Turk & Levoy, 1994) and Objaverse: five meshes spanning 50k–374k vertices, a 750k-vertex subdivided spot, and three full-resolution scan meshes (a gravestone at 0.8M and 2.8M vertices and a jug at 1.0M). Each uses its original texture or, for scans, a baked procedural material.

- **Periodic and high-frequency textures.** Six synthetic materials (stripes and wood at $f \in \{8, 16\}$, Gabor noise at f=16, blobs) at tilings $1\times$, $2\times$, $4\times$, $8\times$ on spot and on a torus (and on the closed slab for the crease study), plus the 20 Objaverse objects with the highest ratio of high-pass to total texture energy among objects passing the UV-fitting filtering rules with at least $1{,}000$ vertices.
- **Text and decals.** Six UV-fitting objects with a text label or logo decal composited onto the texture.
- **Re-discretisations.** Twenty UV-fitting objects, each with QEM simplification to 25% and 10% of the faces, isotropic remeshing at the mean edge length, one Loop subdivision step (Loop, 1987), and 200k area-uniform samples as a point cloud.
- **Non-manifold, triangle-soup and scan inputs.** Five non-manifold Thingi10K meshes, five raw Objaverse triangle soups with unmerged vertices, five real scans (GSO, Downs et al., 2022; DTU, Jensen et al., 2014).
- **Photographs.** Eight objects with photographs and meshes from GSO, NeRF-synthetic (Mildenhall et al., 2020) and DTU for inverse rendering.

E BASELINE ADAPTATIONS

- **HKTex** (Foti et al., 2026) uses the official repository with default configuration ($K = 256$, 7×7 operator grid, $\sigma = 0.05$ biharmonic window, 50 candidates, top 30). HKTex-perpart runs one instance per connected component with the kernel budget split proportionally to area. HKTex-simplify decimates to ≤ 50k vertices with QSlim (Garland & Heckbert, 1997), trains, and maps each kernel to the nearest point of the input mesh with its parameters unchanged. HKTex-K512 doubles the modes and the resident basis. The robust principal frame ("+frame" in Table 2) replaces the per-face mean of vertex principal directions by a sign-aligned, smoothed direction field with an axis-aligned fallback where the principal directions are undefined; it is a one-line configuration change we report to the upstream authors.
- **GT texture, vertex colours** (Yuksel et al., 2010)**.** The texture is downsampled with area averaging until its PNG size matches; vertex colours are assigned on a Loop-subdivided mesh (Loop, 1987) with depth chosen to match storage.
- **Intrinsic Neural Fields** (Koestler et al., 2022)**, Texture Fields + RFF** (Oechsle et al., 2019; Tancik et al., 2020)**.** The HKTex paper's configurations, with hidden width matched to storage.
- **Hash-grid neural texture.** Multi-resolution hash encoding (Müller et al., 2022) of surface positions (16 levels, 2 features per level, finest resolution 2048), two-layer decoder of width 64; the table size 2^T is chosen per object so that the zipped float16 size matches the target. We do not add a surface-aware encoding, so that the baseline is the one a reader would build first; a variant with normals concatenated is reported in Table 13.
- **GStex-on-mesh** (Rong et al., 2025)**, 3D Gabor Splatting-on-mesh** (Watanabe et al., 2025)**, 3DGabSplat-on-mesh** (Zhou et al., 2025)**.** Primitives are 2D Gaussians whose centres are constrained to the mesh surface and whose normals equal the interpolated mesh normal; positions move by projection after each step; opacity, scale, rotation and appearance are learned; GStex uses an 8×8 texel patch per primitive, 3D Gabor Splatting four fixed-direction waves in the splat plane, 3DGabSplat two frequency groups. Rendering is the original tile-based rasteriser; the loss is the same image loss as the other methods on the same views.

F ADDITIONAL TABLES

Table 5: **Ablations** on a fixed random sample of 20 UV-fitting objects and the stripes material at f=16. Entries are view-PSNR changes relative to HeatAtlas (Gabor) (or HeatAtlas (const) where marked); $\pm$ is the standard deviation over three paired seeds.

Component	Variant	Δ UV-fitting (random)	Δ stripes f=16
Coordinates	unfolding → Euclidean chord + projection	-0.23 ± 0.13	-0.54 ± 0.23
Coordinates	unfolding → exact geodesics (10–100× cost) (Mitchell et al., 1987)	-0.04 ± 0.06	-0.11 ± 0.09
Kernel	first-order → zeroth-order (no curvature term)	-0.22 ± 0.10	-0.34 ± 0.14
Kernel	anisotropic → isotropic (const)	-7.91 ± 0.27	-2.78 ± 0.20
Kernel	per-kernel t_i → global t	-0.25 ± 0.14	-2.63 ± 0.23
Kernel	geodesic window σ=0.1 → none	-4.24 ± 0.20	-3.27 ± 0.25
Frame	parallel transport → per-face fixed frame	-0.11 ± 0.09	-2.87 ± 0.23
Support	r_i at 1% tail → 0.1% / 10%	$+0.13 \pm 0.10$ / -1.01 ± 0.15	–
Candidates	per-query → per-face cap (50)	-2.62 ± 0.17	–
Support rebuild	every 10 steps → every step (5.6× slower)	-0.26 ± 0.12	–
Gabor waves	W=4 → 1 / → 2 / → 8	-0.71 ± 0.19 / -0.21 ± 0.12 / -0.12 ± 0.11	-1.10 ± 0.30 / -1.07 ± 0.20 / $+0.03 \pm 0.09$
Gabor frequency	learnable → fixed six-level bank	-0.27 ± 0.14	-1.38 ± 0.20
Gabor phase	geodesic → Euclidean projection	$+0.06 \pm 0.07$	-0.10 ± 0.11
Content	Gabor → constant, 2.5× kernels	-1.24 ± 0.19	-2.65 ± 0.26
Content	oscillation in weight, not colour (additive)	-3.61 ± 0.27	-2.40 ± 0.26
Density control	frequency-aware split → plain split	-0.17 ± 0.11	-0.29 ± 0.15
Barrier	on → off (six material-boundary objects)	-4.87 ± 0.27 in 3-px boundary band	–
Proxy	original mesh → coarse proxy (large meshes)	-0.23 ± 0.14	–

Table 6: **Scalability on large meshes.** Initialisation includes all preprocessing; HKTex (Foti et al., 2026) keeps a 51 kB-per-vertex basis resident. Rows 50k–374k are five different meshes measured with a single seed, so initialisation follows each mesh's kernel supports rather than $|\mathcal{V}|$ alone (HKTex jumps at 175k as well, and the proxied 750k spot initialises in 37 s); the mean row averages these five rows (HKTex-simplify over the three rows with a value). The 750k row is subdivided spot with the proxy; the last row contains three full-mesh scans on an 80 GB GPU (Figure 9). HKTex-simplify (Appendix E) is reported relative to HKTex; on the scan rows, where HKTex has no result, relative to HeatAtlas.

	HKTex		HeatAtlas		ours − HKTex	HKTex-simplify
$\lvert\mathcal{V}\rvert$	init	peak GB	init	peak GB	ΔPSNR	ΔPSNR
50k	3.92 min	2.05	1.55 s	0.07	+0.14	−0.84
100k	11.25 min	4.26	3.27 s	0.21	−0.17	−1.26
175k	56 min	18.1	22 s	0.4	+0.5	−0.57
250k	64 min	25.9	32 s	0.6	0.0	–
374k	74 min	38.7	45 s	0.9	+0.1	–
mean	41.8 min	17.8	20.8 s	0.44	+0.11	−0.89
750k	OOM (35.7 GiB)	–	37 s	0.48	–	–
0.8M / 1.0M / 2.8M	OOM / OOM / OOM	–	73 / 111 / 209 s	2.2 / 3.1 / 8.7	–	−1.06 / −0.58 / −0.57

Table 7: **Structural-failure set by category.** Surface PSNR; Mean is the arithmetic mean over the four equally sized categories. $\pm$ is the standard deviation over three seeds. HeatAtlas (Euclid) uses the same kernels with chord distance and tangent projection instead of unfolding. Bold/underline: best/second best.

Method	Multi-part	Planar	Thin sheet	Near contact	Mean
HKTex (Foti et al., 2026)	19.48 ± 0.26	18.15 ± 0.26	19.65 ± 0.26	19.34 ± 0.28	19.16 ± 0.17
HKTex-perpart (Foti et al., 2026)	29.38 ± 0.21	18.17 ± 0.31	19.54 ± 0.28	19.28 ± 0.29	21.59 ± 0.14
HKTex-simplify (Foti et al., 2026)	20.52 ± 0.22	24.53 ± 0.25	20.18 ± 0.21	21.31 ± 0.25	21.64 ± 0.15
HKTex-perpart+frame (Foti et al., 2026)	28.41 ± 0.15	31.17 ± 0.18	24.35 ± 0.15	22.54 ± 0.16	26.62 ± 0.18
Hash-grid NT (Müller et al., 2022)	28.13 ± 0.20	30.80 ± 0.16	24.55 ± 0.26	23.81 ± 0.27	26.82 ± 0.11
HeatAtlas (Euclid)	$\mathbf{32.36 \pm 0.17}$	$\mathbf{33.24 \pm 0.17}$	22.04 ± 0.29	26.83 ± 0.22	28.62 ± 0.14
HeatAtlas (const)	$\mathbf{32.36 \pm 0.14}$	31.62 ± 0.16	$\mathbf{33.83 \pm 0.18}$	$\underline{31.35 \pm 0.24}$	$\underline{32.29 \pm 0.11}$
HeatAtlas (Gabor)	$\underline{32.05 \pm 0.16}$	$\underline{32.24 \pm 0.15}$	$\underline{33.33 \pm 0.18}$	$\mathbf{31.71 \pm 0.18}$	$\mathbf{32.33 \pm 0.15}$

Table 8: **UV-texture fitting.** Means over 313 objects at matched storage; ΔPSNR is relative to reproduced HKTex. ± is the standard deviation over three seeds. Bold/underline: best/second best, excluding the higher-capacity patch control.

Method	Storage (kB)	PSNR↑	ΔPSNR	SSIM↑	LPIPS↓
GT texture, downsampled	102.79	26.10	-4.97	0.891	0.101
Vertex colours (Yuksel et al., 2010)	103.53	28.39	-2.68	0.885	0.086
Intrinsic Neural Fields (Koestler et al., 2022)	94.37	29.41 ± 0.08	-1.66 ± 0.09	0.916 ± 0.0011	0.081 ± 0.0011
Texture Fields (Oechsle et al., 2019) + RFF (Tancik et al., 2020)	100.90	28.62 ± 0.13	-2.45 ± 0.14	0.938 ± 0.0022	0.060 ± 0.0018
Hash-grid neural texture (Müller et al., 2022)	106.12	$\underline{33.11 \pm 0.08}$	$\underline{+2.04 \pm 0.09}$	0.948 ± 0.0011	$\underline{0.059 \pm 0.0012}$
HKTex (Foti et al., 2026) (reproduced)	96.73	31.07 ± 0.07	0	0.936 ± 0.0011	0.068 ± 0.0012
HKTex-simplify (Foti et al., 2026) ($\leq$ 50k)	102.07	30.86 ± 0.08	-0.21 ± 0.09	0.926 ± 0.0012	0.089 ± 0.0016
HeatAtlas (const)	97.57	32.07 ± 0.07	$+1.00 \pm 0.11$	$\underline{0.950 \pm 0.0012}$	0.063 ± 0.0012
HeatAtlas (Gabor)	101.55	$\mathbf{33.21 \pm 0.11}$	$\mathbf{+2.14 \pm 0.15}$	$\mathbf{0.953 \pm 0.0015}$	$\mathbf{0.043 \pm 0.0015}$
HeatAtlas (patch) (higher-capacity control)	91.60	33.15 ± 0.09	$+2.08 \pm 0.13$	0.964 ± 0.0016	0.038 ± 0.0017

Table 9: **Theory validation.** UV-fitting measurements and analytic surfaces.

Quantity	Measurement	Result
Effective t of trained HeatAtlas kernels	inverted from fitted width	median 0.0029 (spot 0.0028)
Required modes $K^* = \frac{A}{4\pi t} \ln \frac{1}{\varepsilon}$	$\varepsilon = 10^{-2}$	median 437 (spot 418)
HKTex PSNR vs. K (64→1024, 8× tiling)	Fig. 2	256→512: ≈ +1.0; 512→1024: ≈ +0.4 dB
Unwindowed zeroth- / first-order kernel vs. exact, $t = 0.005$	rel. error, sphere ($\rho = 0.25$)	3.6% / 0.1%
Zeroth-order kernel vs. exact, $t = 0.05$	rel. error	15.50%
Curvature correction, PSNR	first − zeroth order	+0.22 dB
Curvature approximation	$\delta_\kappa = \lvert\kappa_D - \kappa\rvert / (\lvert\kappa_D\rvert + \lvert\kappa\rvert + \epsilon)$; PSNR	median 6.42%, p95 20.15%; ΔPSNR +0.04 dB

Table 10: **The learned scale does not follow the window.** (a) HeatAtlas trained with geodesic window width σ on the 20-object UV-fitting sample and stripes at f=16; t_i is the learned per-kernel diffusion time (not inverted from width), and "at cap" is the fraction of kernels at the upper clamp $1/16$. ΔPSNR is view PSNR relative to σ=0.1. (b) HKTex with and without its biharmonic window on 20 textures.

(a) HeatAtlas, window width σ	median t_i	t_i p10–p90	at cap	90% radius (median / p99)	ΔPSNR (UV / stripes)
0.05	0.0058	0.0016–0.022	3.7%	0.043 / 0.098	−0.54 / −0.84
0.1 (default)	0.0052	0.0015–0.020	3.2%	0.056 / 0.110	0 / 0
0.2	0.0057	0.0017–0.025	5.3%	0.069 / 0.16	+0.063 / −0.11
none	0.0064	0.0019–0.029	5.1%	0.13 / 0.19	−4.24 / −3.27
(b) HKTex, biharmonic window	fitted t_{eff}			90% radius (median)	ΔPSNR, zeroed beyond r=0.2
window on (default)	0.0058			0.051	−0.014
window off	0.0051			0.097	−0.048

Table 11: **Multi-view inverse rendering** and real photographs. Training time is end-to-end wall time, including preprocessing, relative to HKTex. ± is the standard deviation over three seeds. Bold/underline: best/second best; storage is matched and therefore not ranked.

Method	Storage (kB)	PSNR↑	LPIPS↓	Train time	Photo PSNR
NVDiffRec (adapted) (Munkberg et al., 2022)	81.39	29.05 ± 0.11	0.067 ± 0.0016	$0.97 \pm 0.03\times$	$\underline{27.21 \pm 0.22}$
HKTex (Foti et al., 2026)	80.06	$\underline{31.28 \pm 0.09}$	$\underline{0.057 \pm 0.0016}$	$1.0 \pm 0.03\times$	26.69 ± 0.21
HeatAtlas (const)	76.18	30.35 ± 0.10	0.063 ± 0.0016	$\mathbf{0.41 \pm 0.02\times}$	26.49 ± 0.20
HeatAtlas (Gabor)	79.90	$\mathbf{31.63 \pm 0.10}$	$\mathbf{0.046 \pm 0.0013}$	$\underline{0.68 \pm 0.03\times}$	$\mathbf{29.95 \pm 0.31}$

Table 12: **Transfer to re-discretisations.** Trained on the original mesh of 20 UV-fitting objects, evaluated on the variant without retraining; Δ is view PSNR relative to the original mesh. $\pm$ is the standard deviation over three paired seeds. For HKTex (Foti et al., 2026), the basis is recomputed on the variant and centres are projected; the hash grid and HeatAtlas are queried directly at the new surface samples. Bold/underline: best/second best.

Variant	HKTex ΔPSNR	Hash grid ΔPSNR	HeatAtlas ΔPSNR
QEM 25% (Garland & Heckbert, 1997)	-5.52 ± 0.26	$\mathbf{-0.39 \pm 0.07}$	$\underline{-0.63 \pm 0.11}$
QEM 10%	-10.44 ± 0.38	$\mathbf{-0.36 \pm 0.08}$	$\underline{-0.82 \pm 0.14}$
Isotropic remesh	-3.75 ± 0.22	$\mathbf{-0.28 \pm 0.06}$	$\underline{-0.33 \pm 0.10}$
Loop subdivision (Loop, 1987)	-7.82 ± 0.27	$\underline{-0.03 \pm 0.05}$	$\mathbf{+0.00 \pm 0.05}$
Point cloud, 200k points	n/a	$\mathbf{-0.49 \pm 0.09}$	$\underline{-1.02 \pm 0.17}$

Table 13: **Non-manifold, triangle soup and scan inputs.** View PSNR. $\pm$ is the standard deviation over three seeds. The cotangent-Laplacian HKTex baseline is not applicable to non-manifold and triangle-soup inputs; "n/a" denotes that no aggregate PSNR is reported. The non-manifold Laplacian of Sharp & Crane (2020b) permits evaluation but does not repair multi-component or duplicate-vertex issues. On non-manifold inputs the transported frame is discontinuous across non-manifold edges, so Gabor phases do not align there, while constant content has no phase. Bold/underline: best/second best.

Method	Non-manifold (5)	Triangle soup (5)	Scans (5)
HKTex (Foti et al., 2026), cotangent Laplacian	n/a	n/a	23.90 ± 0.31
HKTex (Foti et al., 2026), non-manifold Laplacian (Sharp & Crane, 2020b)	26.72 ± 0.27	23.10 ± 0.33	26.36 ± 0.28
Hash-grid neural texture (Müller et al., 2022)	$\underline{32.24 \pm 0.21}$	29.03 ± 0.26	$\mathbf{31.22 \pm 0.24}$
Hash-grid + normals (Müller et al., 2022)	31.79 ± 0.20	$\underline{30.36 \pm 0.25}$	28.62 ± 0.23
HeatAtlas (const)	$\mathbf{32.83 \pm 0.17}$	29.35 ± 0.22	28.37 ± 0.23
HeatAtlas (Gabor)	29.48 ± 0.23	$\mathbf{32.24 \pm 0.23}$	$\underline{31.02 \pm 0.23}$

Table 14: **Efficiency on ten objects drawn from the UV-fitting set.** HKTex and HeatAtlas (const) are means over the ten meshes and three seeds; HeatAtlas (Gabor) uses single-seed measurements. Bold/underline: best/second best.

Metric	HKTex	HeatAtlas (const)	HeatAtlas (Gabor)	Gain (const)
Initialisation (s)	163.8	**4.0**	$\underline{4.64}$	41.0×
Median step (ms)	23.8	**15.6**	$\underline{17.13}$	1.5×
Optimisation, 5,000 steps (s)	**141.8**	$\underline{144.0}$	161.00	–
Peak training memory (GB)	2.95	**0.36**	$\underline{0.42}$	8.2×
Query throughput (M points/s)	0.145	**1.44**	$\underline{1.19}$	10×
Scalars per kernel	8	8	23	–
Resident geometry per vertex	51 kB	$O(1)$ per proxy face	$O(1)$	–

Table 15: **Crease study on closed slabs** (seed 0, view PSNR, storage matched). The planar top with a 90° edge is where geodesic and Euclidean phase first differ measurably.

Group	#	HeatAtlas (const)	HeatAtlas (const), 2.5×	HeatAtlas (Gabor)	HeatAtlas (Gabor), Euclid. phase
slab (planar top)	8	28.79	31.54	**35.01**	34.86
wavy slab	8	29.12	31.80	**34.53**	34.36
torus	8	31.74	33.72	35.59	**35.60**
spot	6	33.23	34.21	35.28	**35.30**
all synthetic	30	30.55	32.72	**35.09**	35.01

Table 16: **Tiling study on spot** (single seed, surface PSNR): the same texture tiled 1×–8× with 5,000 kernels, and the effect of kernel budget and per-kernel scale at 8×. For HKTex, $+t_i$ adds a learnable per-kernel diffusion time to the spectral evaluator ($e^{-\lambda_k t_i}$ per kernel) at unchanged K.

	1×	2×	4×	8×	8×, 20k	8×, 20k + t_i
HKTex (Foti et al., 2026)	33.7	29.3	24.7	20.6	21.1	22.4
HeatAtlas (const)	35.1	30.7	26.0	21.9	22.7	**24.5**

G ADDITIONAL DIAGNOSTICS

The following measurements isolate geometric and numerical properties that complement the main tables.

- **Kernel radius distribution.** Across 623,116 trained HeatAtlas kernels on 130 UV-fitting objects, the 90%-contribution radius has median 0.056 and 99th percentile 0.110 (Figure 2). On spot the corresponding values are 0.067 and 0.163.
- **Unfolding accuracy.** Against MMP exact geodesics on spot at radius 0.4, unfolded distance has median relative error 0.03% and p99 6.9%.
- **Analytic curvature validation.** On a sphere of radius 0.25 at $t = 0.005$, the zeroth- and first-order kernels have relative errors 3.6% and 0.1%, respectively (Figure 2, Table 9).
- K **sweep.** On spot, the predicted $K^* = 418$ lies at the observed knee of the PSNR curve; increasing K from 256 to 512 gains 1.0 dB and 512 to 1024 gains 0.4 dB (Figure 2).
- **Candidate coverage.** On a 121-vertex mesh, a per-face cap leaves 36.5% of the surface with $\sum_j f_j < 1$, whereas per-query selection reduces that fraction to 1.2%.

H ADDITIONAL QUALITATIVE FIGURES

Figures 7–10 show transfer, inverse rendering, failure cases, large meshes, edits and barriers.

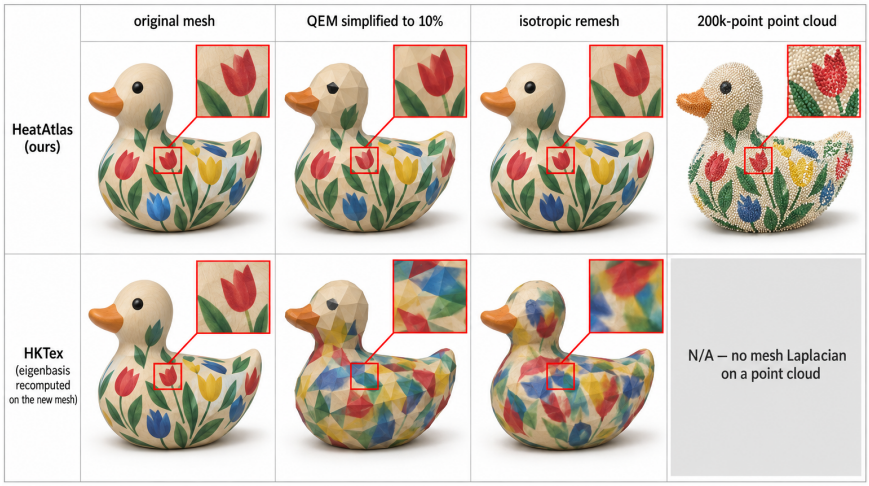


Figure 7: **Transfer without retraining.** A texture trained on the original mesh, evaluated on a QEM 10% simplification, an isotropic remesh and a 200k-point cloud. HKTex with an eigenbasis recomputed on the new mesh loses the pattern; a point cloud has no mesh Laplacian.



Figure 8: **Inverse rendering and failure cases.** Top: photographs vs. HKTex, HeatAtlas and NVDiffRec (Munkberg et al., 2022). Bottom: (a) text under Gabor content, (b) a rim thinner than the kernel support, where supports wrap around the edge, (c) a frame jump on a planar region with a sharp crease, (d) a thin stem not seeded by error-driven splitting.

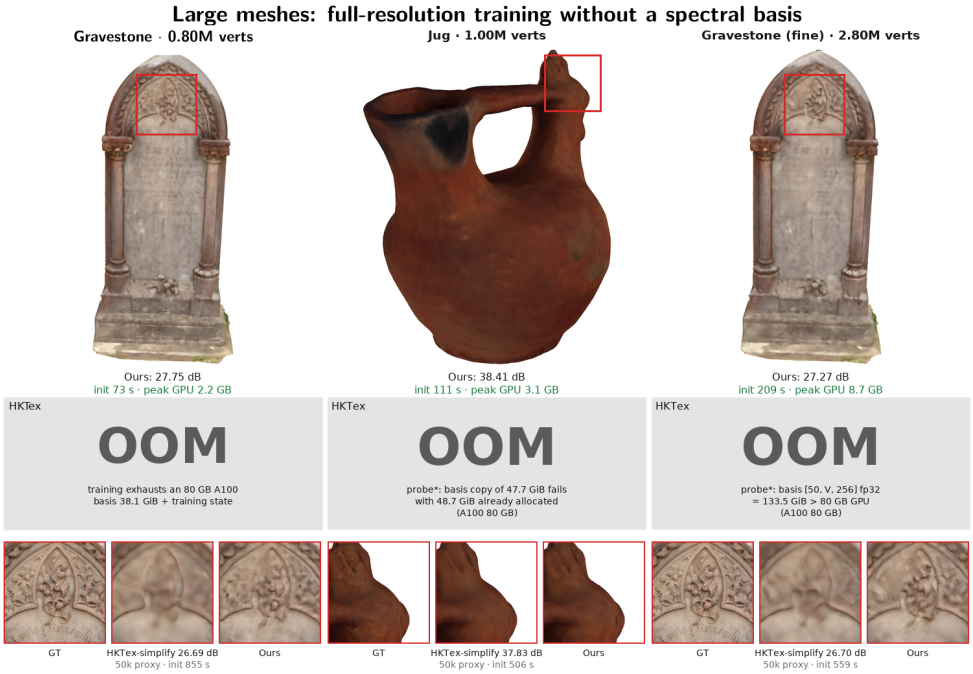


Figure 9: **Large meshes.** HeatAtlas trains on the full 0.80M-, 1.00M- and 2.80M-vertex scans (view PSNR 27.75, 38.41 and 27.27 dB; initialisation 73, 111 and 209 s; peak GPU memory 2.2, 3.1 and 8.7 GB). HKTex exhausts an 80 GB A100 during full-resolution training on all three scans. Insets compare GT, HKTex-simplify on a 50k proxy (26.69, 37.83 and 26.70 dB) and ours.

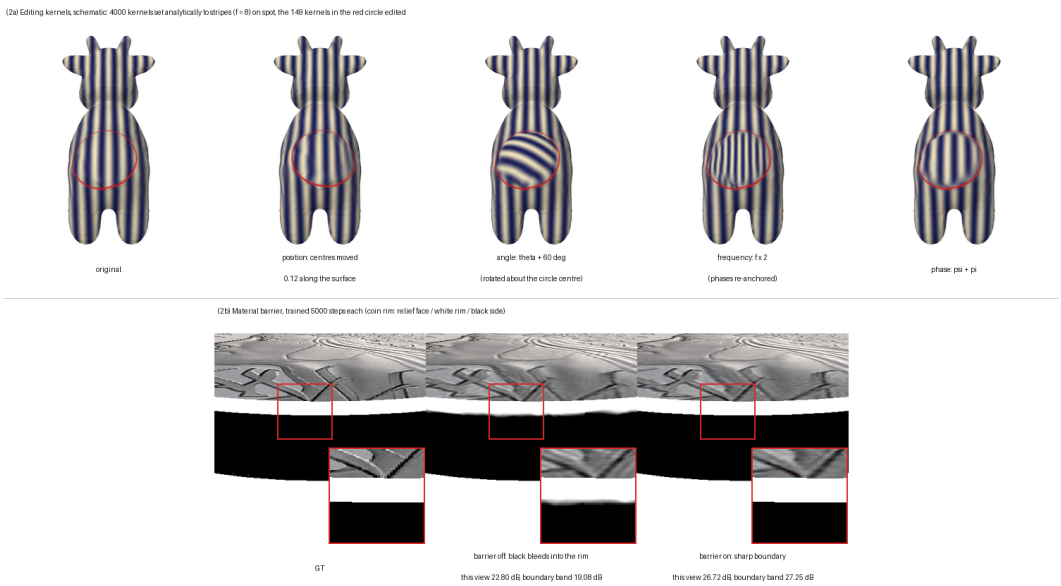


Figure 10: **Local edits and barriers.** Top: $4{,}000$ kernels set analytically to stripes ($f = 8$) on spot; the 149 kernels inside the red circle are edited by moving their centres 0.12 along the surface, rotating θ by $60°$, doubling the frequency, and shifting the phase by π. Bottom: a coin rim with relief face, white rim and black side, trained for $5{,}000$ steps. Without a barrier black bleeds into the rim (view 22.80 dB, boundary band 19.08 dB); with the barrier the boundary stays sharp (26.72 and 27.25 dB).

I Discussion and limitations

Locality and intrinsicness. Locality drives the gains in scale, robustness and transfer, because no step depends on a global eigensystem. Intrinsic charts matter where Euclidean neighbourhoods are wrong, across thin sheets, contacts and creases, and they supply the transported frame in which directional content defines its frequencies and phases.

Explicit primitives and neural textures. Below 100 kB HeatAtlas (Gabor) gives the highest PSNR on the evaluated UV-fitting subset, and on the two periodic materials it leads at every budget; at 200–400 kB the hash grid attains higher PSNR on that UV-fitting subset (Figure 3). The two offer different properties: HeatAtlas is a few thousand explicit, seam-free primitives with sharp barriers, local edits and a non-negative normalised composition queried directly by a ray tracer, whereas a hash grid is a volumetric raster in which nearby sheets share cells and no element can be selected or moved.

Coverage of isolated regions. Density control is error-driven, so splitting refines regions that already carry kernels but cannot seed a region that none reaches. A thin stem separated from the rest of the surface stays under-covered (Figure 8d). Coverage-driven insertion, which places new kernels where $\sum_j f_j$ is small, directly extends the coverage-normalised composition.

Content families. Intrinsic Gabor content represents locally periodic and oriented appearance; on non-periodic textures it behaves like a smooth intra-kernel colour gradient and does not represent text or logos (Figure 8a). The learnable patch covers these cases at the price of turning primitives into neural tiles with a shared decoder.

Order of the expansion. The kernel equation 5 is first-order in t and second-order in d. Trained kernels have median $t_i = 0.0052$ (Table 10); at $t = 0.005$ the first-order error is 0.1%, while kernels at $t \geq 0.05$ would need higher-order terms of equation 2. Using surface curvature in place of anisotropic-metric curvature changes PSNR by only $+0.04$ dB (Table 9).

Frames and barriers. The unfolded frame is a discrete parallel transport and is consistent only up to holonomy, so Gabor phases spanning a high-curvature cone vertex are ambiguous (Figure 8c

shows a related frame jump where a planar region meets a sharp crease). Barriers follow mesh edges, so a boundary that cuts across faces is approximated by an edge path.

Intrinsic accuracy on smooth objects. On smooth, isolated objects a Euclidean projection is within 0.23 dB of unfolded coordinates, and geodesic and Euclidean Gabor phase agree within 0.1 dB across the displayed torus and spot conditions. The benefit of intrinsic charts concentrates on thin sheets, contacts and creases.

Scale and tuning. The proxy removes tessellation from support cost but not input size from I/O and rasterisation. All learning rates and density-control settings are inherited from HKTex and were not re-tuned for HeatAtlas.