UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

Dipartimento di Scienze Fisiche, Informatiche e Matematiche

Master's Degree Course in Mathematics (Data Science)

ON SMOOTH ATLASES OVER TRIANGLE MESHES: AN EMPIRICAL STUDY OF DALa FOR SMOOTH GEOMETRY PROCESSING

Supervisors
Prof. Fabio Pellacini
Prof. Claudio Mancinelli

Candidate
Matteo Lugli

Academic Year 2025/2026

This first page is dedicated to some of my favorite quotes, which reflect the spirit of this degree course, this work and the way I approached it. The interpretation of each one of them is up to the reader and I hope he/she/they will enjoy them as much as I do.

«*When a man with a .45 meets a man with a rifle,*
the man with the pistol is a dead man.»
THE MAN WITH NO NAME, *A Fistful of Dollars* (1964)

«*Could be worse... could be raining.*»
IGOR, *Young Frankenstein* (1974)

«*I've seen things you people wouldn't believe...*»
ROY BATTY, *Blade Runner* (1982)

«*This is an imaginary story... aren't they all?*»
ALAN MOORE, *Whatever Happened to the Man of Tomorrow?* (1986)

«*It can't rain all the time.*»
ERIC DRAVEN, *The Crow* (1994)

«*I teach this ****, I didn't say I know how to do it.*»
SEAN MCGUIRE, *Good Will Hunting* (1997)

«*The Dude abides.*»
JEFFREY THE DUDE LEBOWSKI, *The Big Lebowski* (1998)

«*We accept the reality of the world with which we're presented.*»
CHRISTOF, *The Truman Show* (1998)

«*Sand is overrated. It's just tiny, little rocks.*»
JOEL BARISH, *Eternal Sunshine of the Spotless Mind* (2004)

«*Don't look for gods in men, you'll always be disappointed.*»
ODYSSEUS, *The Odyssey* (2026)

For my parents.
For my sister.
For Sabrina.
For myself.

Contents

List of Figures

List of Tables

Introduction

A triangle mesh is a piecewise-linear object. It stores a finite set of vertex positions and declares the surface to be flat in between, which is enough to represent shape convincingly, but does not carry a globally smooth differential structure across its edges. Face normals are well defined, but curvature and the classical differential operators are not pointwise defined on the entire piecewise-linear surface. To this difficulty — the object one computes on is not the object the theory is written for — there are two possible responses, and the contrast between them is the setting of this thesis.

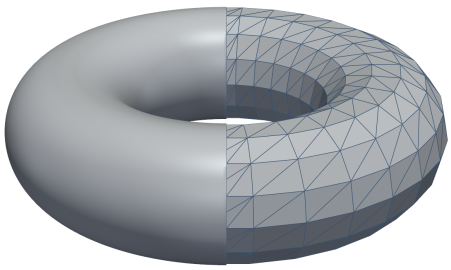

Figure 1. The same torus, seen as the smooth surface it is meant to represent (left) and as the triangle mesh that represents it (right).

The first is to redefine the quantities on the mesh itself and discrete differential geometry does exactly this: it takes identities that hold in the smooth setting and promotes them to definitions in the discrete one, giving the angle defect for the Gaussian curvature[1] and the cotangent formula for the Laplacian. These estimators are efficient, remarkably accurate on well-shaped meshes and, by construction, tied to the quality of the tessellation, which is rarely something one controls.

The second response reverses the move: instead of redefining the differential operators on a piecewise-linear domain, it reconstructs a smooth surrogate. One builds a

[1]The *angle defect* at a vertex is $2\pi - \sum_t \theta_t$, the amount by which the angles θ_t of the incident triangles fail to close a flat disc; divided by the area attached to the vertex it estimates the Gaussian curvature there. It is the discrete baseline against which every experiment of Chapter 7 is measured (Section 7.1), and the cotangent formula is derived in Section 3.2.

differentiable parametrisation *over* the mesh — an atlas of local charts blended into a globally compatible map — and reads the classical definitions off that. The appeal is that nothing needs reinterpreting: the Gaussian curvature is the one given by the first and second fundamental forms, evaluated at any point of any triangle and not only at the vertices. The idea is not new, but it has recently become practical: [18] construct a G^1 interpolant by blending vertex-centric local geometries, and [19] (Chapter 4) extends it to prescribed C^k regularity while making the whole evaluation pipeline differentiable, so that quantities of arbitrary order are evaluated by automatic differentiation, without finite differences.

There is a subtlety in this construction, and it motivates the present work. Such an atlas is fitted to the mesh using *order-zero* information only: the vertex positions — no constraint is imposed on normals, on curvature, or on any higher-order quantity. Positional accuracy is therefore optimised directly, whereas the same reconstruction is then queried for a *second-order* quantity. Whether the second answer inherits the accuracy of the first is not obvious, and it is the question this thesis addresses:

does a positionally faithful C^k reconstruction of a triangle mesh yield a faithful Gaussian curvature?

The question is posed empirically, on a testbed designed so that the answer can be measured rather than argued. The torus of revolution has a curvature known in closed form[2] and, unlike the sphere, is not the surface that the bending energy of the fit already prefers; it is tessellated in four qualitatively different ways at three resolutions, and six estimators — the discrete baseline, an oracle driven by the analytic target, and four variational fits — are compared against the exact value. On this common ground a sequence of controlled experiments varies one ingredient at a time: the tessellation, the interpolation constraint, the target of the attractor, the quadrature that assembles the fit, the blending weights, the regularisation of the fit, and the polynomial degree of the charts. A final experiment carries the atlas into a downstream application (functional maps) to ask whether smooth off-vertex evaluation helps a pipeline that consumes scalar fields rather than derivatives.

The evidence supports a qualified answer, which separates into a positive and a negative finding. On the positive side, selected atlas reconstructions show a decreasing vertex-curvature error over the tested refinements, including on the irregular tessellations where the angle-defect error does not decrease. This is

[2]For the torus of major radius R and minor radius r, parametrised by u around the tube and v around the axis, the Gaussian curvature is $K = \cos u \,/\, [\, r\,(R + r\cos u)\,] = (\rho - R)/(r^2\rho)$ with $\rho = \sqrt{x^2 + y^2}$. The right-hand form evaluates it from Cartesian coordinates alone. At reconstructed points that lie slightly off the analytic surface, it is used as a consistent radial extension of the reference value; see Section 7.1.

qualitatively consistent with the robustness of fitting-based estimators [16], although the available theorem does not apply directly to DALa's global variational fit. On the negative side, the experiments show that positional accuracy is a poor predictor of differential accuracy: across the sampled triangle interiors, median positional errors are of order 10^{-5} while curvature errors are three orders of magnitude larger, and the two quantities are only loosely correlated. Tracking that gap occupies the second half of the thesis. The exact product-rule decomposition identifies additional terms introduced by the partition of unity; three controlled experiments support this interpretation, and a regularisation that keeps neighbouring charts in agreement where the weights vary fastest substantially reduces the measured error in the tested configurations.

Structure of the thesis. Part I collects the background: the differential geometry of surfaces (Chapter 1), the spectral theory attached to the Laplace–Beltrami operator (Chapter 2) and its discretisation (Chapter 3), the smooth atlas that is the object of study (Chapter 4), and functional maps (Chapter 5), the application used at the end. Part II opens with the motivation and the claims (Chapter 6), reports the experiments (Chapter 7) and interprets them (Chapter 8), where the theoretical reading of the results is developed. Part III draws the conclusions (Chapter 9) and outlines the work they open (Chapter 10). Two appendices recall the point-set topology and the matrix algebra used throughout.
The reconstruction library is externally supplied and the DALa language and atlas construction are described in a manuscript in preparation [19]; everything around it — the test surfaces, the drivers, the error measurements, the figures and the renders — was written for this work, so that every number reported here is reproducible with a single command.

Notation and Conventions

This section fixes the conventions and collects the symbols used throughout the thesis, so that it may serve the reader as a quick reference. Symbols attached to a single construction are instead defined where they first appear — not listed here.

Smoothness. The term "smooth" means C^∞; when a finite order is intended it is stated explicitly as C^k. The word "differentiable" means differentiable to the order required by the operation under discussion. When no ambiguity arises we write 0 for the origin of $\mathbb{R}^n$.

Smooth vs. discrete surface. The smooth Riemannian surface is denoted M and its triangle-mesh discretisation $\mathcal{M}$. A symbol introduced in the smooth setting may be reused for an explicitly defined discrete analogue; this convention does not endow the piecewise-linear mesh with smooth objects that have not been separately constructed.

Points, vectors and matrices. Points, scalar functions and their sampled values are set in italic (p, f); ambient vectors in bold ($\mathbf{x}$). Geometric operators and basis matrices are upright capitals (M, W, L, Φ, Λ), while the functional-maps literature conventionally uses bold capitals ($\mathbf{A}, \mathbf{B}, \mathbf{C}$). Spectral coefficient vectors are written a, b and their entries a_i, b_i. The notation $A^\top$ denotes the transpose, I the identity, and $\mathrm{diag}(\cdot)$ a diagonal matrix. Following common usage in geometry processing, M denotes both the smooth manifold and the discrete mass matrix, disambiguated by context.

Reused letters. The letter K is reserved throughout for Gaussian curvature. Subscripts only specify the surface or the role in a comparison: K_S is the curvature of the reconstructed surface, K_{true} the analytic reference and K_{est} a generic estimate; they do not denote different notions of curvature. The lowercase k denotes a regularity order in C^k and, in spectral formulas, the local truncation order. The letter S is a regular embedded surface in Chapter 1 and, from Chapter 4 onwards, the blended surface produced by the atlas; S_T is its expression in the triangle chart U_T. DALa's implementation continuity parameter is written k_{cont} to distinguish it from both K and the local index k.

Einstein summation. An index appearing once as a superscript and once as a subscript in the same term is implicitly summed over its range, the symbol $\sum$ being omitted: thus $X = X^i \partial_i$ stands for $\sum_i X^i \partial_i$, and $\Gamma^k_{ij} X^i Y^j$ is summed over i and j. Superscripts denote contravariant (vector) components, X^i; subscripts denote covariant indices or basis fields, ∂_i and g_{ij}.

Laplacian sign. The Laplace–Beltrami operator is taken positive semidefinite, $\Delta = -\operatorname{div} \nabla$, so that its eigenvalues are non-negative and indexed from zero, $0 = \lambda_0 \leq \lambda_1 \leq \cdots$.

In the following, the reader can find a summary of the main symbols used in the thesis, grouped by context.

Continuous setting

$M,\ (M, g)$	Closed Riemannian surface (manifold)
S	Regular surface embedded in $\mathbb{R}^3$
$T_pM,\ T_pS$	Tangent space at the point p
$\{\partial_i\} = \{\mathbf{x}_u, \mathbf{x}_v\}$	Basis of the tangent space in local coordinates (u^1, u^2)
$\mathbf{N}$	Unit normal field / Gauss map
$\langle \cdot, \cdot \rangle$	Euclidean inner product of $\mathbb{R}^3$
$\langle \cdot, \cdot \rangle_g$	Riemannian (metric) inner product on T_pM
$g,\ g_{ij},\ (E, F, G)$	Riemannian metric / first fundamental form
$\mathrm{d}\mu$	Riemannian area measure, $\sqrt{\det g}\,\mathrm{d}u^1 \mathrm{d}u^2$
K	Gaussian curvature

$\mathfrak{X}(M)$	Smooth vector fields on M
$C^\infty(M)$	Smooth real-valued functions on M
Γ_{ij}^k	Christoffel symbols of the first fundamental form
$\gamma,\ \gamma_v$	Geodesic; geodesic with $\gamma_v(0) = p$, $\gamma_v'(0) = v$
$\exp_p,\ \log_p$	Exponential and logarithmic maps at p
$d(p, q)$	Geodesic distance between two points

Spectral setting

$L^2(M)$	Space of square-integrable functions on M
$\langle\cdot,\cdot\rangle_{L^2}$	L^2 inner product, $\int_M f\, h\, \mathrm{d}\mu$
$\nabla f,\ \operatorname{div} X$	Riemannian gradient of f; divergence of X
Δ	Laplace–Beltrami operator
λ_k	k-th eigenvalue
ϕ_k	k-th eigenfunction, $\Delta\phi_k = \lambda_k\phi_k$
$\hat{f}_k$	Generalized Fourier coefficient $\langle f, \phi_k\rangle_{L^2}$
k	Spectral truncation order
$k_t(x, y)$	Heat kernel at diffusion time t
$\mathrm{hks}(x)$	Heat Kernel Signature of the point x

Discrete setting

$\mathcal{M} = (V, F)$	Triangle mesh with vertex set V and face set F
$v_i,\ n$	Vertex of the mesh; number of vertices $n = \lvert V\rvert$
t_{ijk}	Triangle with vertices v_i, v_j, v_k
$i \sim j$	Adjacency: v_i and v_j share an edge
$\alpha_{ij},\ \beta_{ij}$	Angles opposite the edge (i, j) (cotangent weights)
M	Mass matrix, $\langle f, h\rangle_{L^2} \approx f^\top M h$
W	Cotangent stiffness matrix
$L = M^{-1}W$	Discrete Laplace–Beltrami operator
Φ	Matrix of discrete eigenvectors, $\Phi^\top M \Phi = I$
Λ	Diagonal matrix of the eigenvalues λ_k
ψ_i	Piecewise-linear hat function associated with v_i
$\beta_a, \beta_b, \beta_c$	Barycentric coordinates inside a face
$x_i,\ \omega_i,\ \omega,\ N$	Surface samples, quadrature (area) weights, their vector $\omega = (\omega_1, \ldots, \omega_N)^\top$, and sample count
$\Phi(X)$	Eigenfunctions evaluated at the sample points X

Smooth atlas

(U_i, φ_i)	Chart: open set of the manifold and its coordinate map
$\varphi_j \circ \varphi_i^{-1}$	Transition function between two overlapping charts
C^k	Order of continuity of the reconstructed surface
$S,\ S_T,\ S(p)$	Blended reconstructed surface, its expression on U_T, and its evaluation at p

$S_i,\ \widehat{S}_i$	Vertex polynomial and its pull-back to triangle coordinates, $\widehat{S}_i = S_i \circ \varphi_{T,i}$
$\varphi_{T,i}$	Transition map from triangle-chart to vertex-chart coordinates
$S_T^\star$	Ground-truth surface expressed in the triangle coordinates of U_T
w_i	Blending weight, $\sum_i w_i \equiv 1$ (partition of unity)
e_i	Pulled-back chart disagreement, $e_i = \widehat{S}_i - S_T^\star$
ε	Decay parameter of the blending gate
$K_S(p)$	Gaussian curvature of the reconstructed surface at $S(p)$
k_{cont}	Continuity parameter used by the DALa implementation

Experiments

$R,\ r$	Major and minor radius of the test torus ($R = 1$, $r = 0.4$)
$K_{\text{est}},\ K_{\text{true}}$	Estimated and analytic Gaussian curvature
MAE	Mean absolute error against K_{true}
maxbary	Largest barycentric coordinate of a sample (1 at vertex, 1/3 at centroid)
h	Representative mesh scale in the convergence heuristics

Functional maps

$T : \mathcal{M} \to \mathcal{N}$	Point-to-point map between two shapes
$T^\star$	Ground-truth point-to-point correspondence
T_F	Functional representation (pull-back) of the map T
$\mathbf{C}$	Functional map matrix, $a = \mathbf{C}\, b$
$a,\ b$	Coefficient vectors of functions in the truncated eigenbasis
$\varepsilon_{\text{geo}}(v)$	Normalized geodesic error of the recovered correspondence at v

PART I

Preliminaries

1. Differential Geometry

This chapter collects the differential-geometric tools that the rest of the thesis relies on. The goal is not to provide an exhaustive treatment, but rather a self-contained immersion for the reader to follow the constructions in the next chapters.
The presentation follows the classical references [2] and [1] and all the basic topology used throughout is recalled in Appendix A.

1.1 Regular surfaces, tangent plane and normal

Definition 1.1 (Regular surface)**.** A subset $S \subseteq \mathbb{R}^3$ is a *regular surface* if for every $p \in S$ there exist a neighbourhood $V \subseteq \mathbb{R}^3$ of p, an open set $U \subseteq \mathbb{R}^2$, and a map $\mathbf{x} : U \to V \cap S$ (called a *parametrization*) such that

(i) $\mathbf{x}$ is smooth (C^∞);

(ii) $\mathbf{x} : U \to V \cap S$ is a homeomorphism;

(iii) for every $q \in U$ the differential $d\mathbf{x}_q : \mathbb{R}^2 \to \mathbb{R}^3$ is injective, equivalently $\mathbf{x}_u(q) \times \mathbf{x}_v(q) \neq 0$, where $\mathbf{x}_u = \partial\mathbf{x}/\partial u$ and $\mathbf{x}_v = \partial\mathbf{x}/\partial v$.

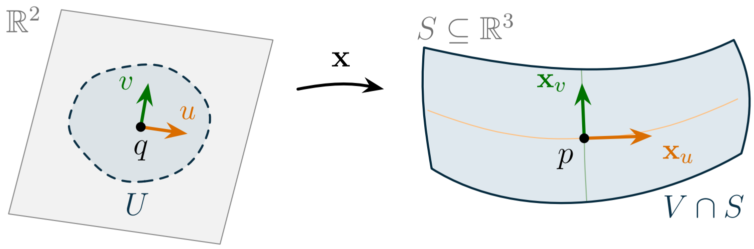

Figure 2. A regular surface is one that looks, near each of its points, like a deformed piece of plane. The parametrization $\mathbf{x}$ maps the open set $U \subseteq \mathbb{R}^2$ homeomorphically onto the patch $V \cap S$, carrying the coordinate lines through q to the coordinate curves through $p = \mathbf{x}(q)$. Condition (iii) requires their velocities $\mathbf{x}_u, \mathbf{x}_v$ to stay linearly independent, which is what prevents the patch from degenerating to a curve or developing a corner.

Proposition 1.2 (Change of parameters)**.** If $\mathbf{x} : U \to S$ and $\bar{\mathbf{x}} : \bar{U} \to S$ are two parametrizations with $W = \mathbf{x}(U) \cap \bar{\mathbf{x}}(\bar{U}) \neq \emptyset$, then the change of coordinates $h = \bar{\mathbf{x}}^{-1} \circ \mathbf{x}$ is a smooth diffeomorphism between open sets of $\mathbb{R}^2$.

Proof. Proving that h is a homeomorphism is immediate: it's a composition of homeomorphisms by condition (ii) of the definition. The same cannot be said for its *smoothness*, which doesn't follow from the definition directly, because $\bar{\mathbf{x}}^{-1}$ is a priori only continuous.
Fix $q \in \mathbf{x}^{-1}(W)$ and let $\bar{q} = h(q)$, so that $\bar{\mathbf{x}}(\bar{q}) = \mathbf{x}(q) =: p$ and write $\bar{\mathbf{x}}(\bar{u}, \bar{v}) = (x(\bar{u}, \bar{v}),\ y(\bar{u}, \bar{v}),\ z(\bar{u}, \bar{v}))$. By condition (iii) the differential $\mathrm{d}\bar{\mathbf{x}}_{\bar{q}}$ is injective, so its 3×2 Jacobian has rank two: some 2×2 minor is non-zero, and after renaming the coordinates of $\mathbb{R}^3$ we may assume

$$\frac{\partial(x, y)}{\partial(\bar{u}, \bar{v})}(\bar{q}) \neq 0.$$

Extend $\bar{\mathbf{x}}$ to the map

$$\begin{array}{rccc} F: & \bar{U} \times \mathbb{R} & \longrightarrow & \mathbb{R}^3 \\ & (\bar{u}, \bar{v}, t) & \longmapsto & \big(x(\bar{u}, \bar{v}),\ y(\bar{u}, \bar{v}),\ z(\bar{u}, \bar{v}) + t\big), \end{array}$$

which is smooth and satisfies $F(\bar{u}, \bar{v}, 0) = \bar{\mathbf{x}}(\bar{u}, \bar{v})$. Its Jacobian determinant at $(\bar{q}, 0)$ is exactly the minor above, hence non-zero, so by the inverse function theorem (Theorem A.6) F is a diffeomorphism of a neighbourhood of $(\bar{q}, 0)$ onto a neighbourhood of p in $\mathbb{R}^3$; in particular F^{-1} is smooth there.
Now restrict to the surface: by continuity of $\mathbf{x}$ there is a neighbourhood V of q with $\mathbf{x}(V)$ contained in that neighbourhood of p, and on V

$$h = \bar{\mathbf{x}}^{-1} \circ \mathbf{x} = \pi \circ F^{-1} \circ \mathbf{x},$$

where $\pi(\bar{u}, \bar{v}, t) = (\bar{u}, \bar{v})$: indeed F^{-1} agrees with $\bar{\mathbf{x}}^{-1}$ on points of the surface, up to the third coordinate $t = 0$ that π discards. The right-hand side is a composition of smooth maps, so h is smooth at q. As q was arbitrary, h is smooth on $\mathbf{x}^{-1}(W)$; exchanging the roles of $\mathbf{x}$ and $\bar{\mathbf{x}}$ shows the same for h^{-1}. □

Observation 1.3. This proposition is what makes every smooth notion on S well defined, independent of the chosen parametrization. It is also the seed of the abstract notion of manifold formalized in Section 1.2: a regular surface is precisely a two-dimensional differentiable manifold embedded in $\mathbb{R}^3$.

Definition 1.4 (Tangent plane)**.** Let $\mathbf{x} : U \to S$ be a parametrization with $\mathbf{x}(q) = p$. The *tangent plane* to S at p is the two-dimensional vector subspace

$$T_pS = \mathrm{d}\mathbf{x}_q(\mathbb{R}^2) = \mathrm{span}\{\mathbf{x}_u(q), \mathbf{x}_v(q)\} \subseteq \mathbb{R}^3. \tag{1.1}$$

Equivalently, $T_pS = \{\alpha'(0) : \alpha \text{ smooth curve on } S,\ \alpha(0) = p\}$, which does not depend on the chosen parametrization.

Definition 1.5 (Gauss map and unit normal)**.** A regular surface S is *orientable* if it admits a smooth unit normal field $\mathbf{N}$ defined on all of S; locally

$$\mathbf{N} = \frac{\mathbf{x}_u \times \mathbf{x}_v}{\|\mathbf{x}_u \times \mathbf{x}_v\|}. \tag{1.2}$$

The *Gauss map* is $\mathbf{N} : S \to S^2$, valued in the unit sphere.

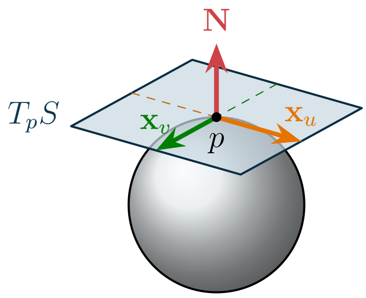

Figure 3. The ordered tangent basis $(\mathbf{x}_u, \mathbf{x}_v)$ spans T_pS and fixes its orientation: the right-hand rule gives the displayed unit normal $\mathbf{N}$.

Definition 1.6 (Differential of a map between surfaces)**.** Let $\varphi : S_1 \to S_2$ be smooth. For $v \in T_pS_1$, choosing a curve α on S_1 with $\alpha(0) = p$, $\alpha'(0) = v$, one sets

$$\mathrm{d}\varphi_p : T_pS_1 \to T_{\varphi(p)}S_2, \qquad \mathrm{d}\varphi_p(v) = \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} (\varphi \circ \alpha), \tag{1.3}$$

which is well defined and linear.

Remark 1.7. The tangent plane T_pS is the linear (flat) approximation of S near p, and $\mathbf{N}$ is its normal direction. The manifold-based pipeline of the thesis flattens each *vertex star* onto such a plane, where the familiar Euclidean tools are available.

1.2 Differentiable manifolds

Surfaces in $\mathbb{R}^3$ are described from the outside, through an ambient space. The manifold viewpoint removes the ambient space and describes the geometry intrinsically: this is the language in which we will set up the problem in the following chapters.

Definition 1.8 (Chart)**.** Let M be a topological space (Hausdorff, second countable). An n-dimensional *chart* is a pair (U_i, φ_i) with $U_i \subseteq M$ and $\widetilde{U}_i \subseteq \mathbb{R}^n$ open and $\varphi_i : U_i \to \widetilde{U}_i$ a homeomorphism. The components of φ_i are the *local coordinates*.

Definition 1.9 (Transition function and C^k-compatibility)**.** Given two charts (U_i, φ_i) and (U_j, φ_j) with $V := U_i \cap U_j \neq \emptyset$, the *transition function* is

$$\varphi_j \circ \varphi_i^{-1} : \varphi_i(V) \subseteq \mathbb{R}^n \longrightarrow \varphi_j(V) \subseteq \mathbb{R}^n. \tag{1.4}$$

The charts are C^k*-compatible* if both this map and its inverse are of class C^k.

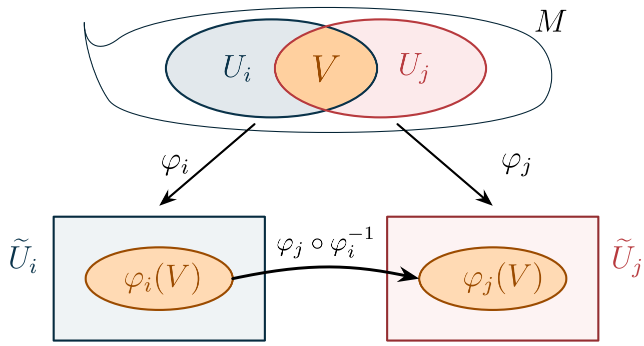

Figure 4. Two overlapping charts on M. The transition function maps the highlighted set $\varphi_i(V)$ to $\varphi_j(V)$, where $V = U_i \cap U_j$.

Remark 1.10. A transition function is a map between open sets of $\mathbb{R}^n$, where ordinary calculus applies: requiring it to be C^k is a concrete, verifiable condition, not an abstraction. The smoothness of M is encoded entirely in the transition functions; a single chart, being merely a homeomorphism, carries only topological (C^0) information.

Definition 1.11 (Atlas and differentiable manifold)**.** An *atlas* of class C^k on M is a family of charts $\mathcal{A} = \{(U_i, \varphi_i)\}_i$ such that the U_i cover M and any two charts are pairwise C^k-compatible. A *differentiable manifold* of dimension n is a space M (Hausdorff, second countable) equipped with a C^∞ atlas of n-dimensional charts.

Observation 1.12. The same M admits many atlases; two atlases are *equivalent* if their union is again an atlas. Equivalent atlases define the same notion of smoothness on M. Taking the union of all equivalent atlases yields the *maximal atlas*, the *differentiable structure*: a single, well-defined object, independent of the particular pages one happened to draw first. In practice it suffices to exhibit any one atlas to fix it.

Example 1.13. The sphere $S^2 = \{(x, y, z) : x^2 + y^2 + z^2 = 1\}$ cannot be covered by a single chart, but two stereographic projections suffice:

$$\varphi_N(x, y, z) = \left(\tfrac{x}{1-z}, \tfrac{y}{1-z}\right), \qquad \text{from the north pole } N = (0, 0, 1), \tag{1.5}$$

$$\varphi_S(x, y, z) = \left(\tfrac{x}{1+z}, \tfrac{y}{1+z}\right), \qquad \text{from the south pole } S = (0, 0, -1). \tag{1.6}$$

Each chart covers the sphere minus one point; together they cover S^2. On the overlap, which in coordinates is $\mathbb{R}^2 \setminus \{0\}$, the transition function is the inversion $u \mapsto u/\|u\|^2$, manifestly C^∞. Hence the two charts give S^2 the structure of a differentiable 2-manifold.

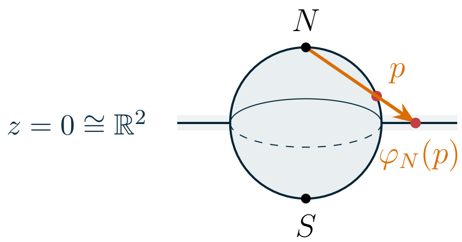

Figure 5. The line through N and p meets the plane at $\varphi_N(p)$. A second projection from S covers the missing pole N.

Definition 1.14 (Smooth functions and maps)**.** A function $f : M \to \mathbb{R}$ is C^k if, for every chart (U_i, φ_i), the coordinate representation $f \circ \varphi_i^{-1} : \widetilde{U}_i \to \mathbb{R}$ is C^k. A map $F : M \to N$ between manifolds is C^k if $\psi \circ F \circ \varphi^{-1}$ is, for every pair of charts φ on M and ψ on N. The C^k-compatibility of the transitions guarantees that these definitions do not depend on the chosen charts.

Definition 1.15 (Tangent space (intrinsic))**.** Fix $p \in M$. On the set of smooth curves $\alpha : (-\varepsilon, \varepsilon) \to M$ with $\alpha(0) = p$ introduce the equivalence relation

$$\alpha \sim \beta \iff (\varphi \circ \alpha)'(0) = (\varphi \circ \beta)'(0) \text{ for one (hence every) chart } \varphi. \tag{1.7}$$

The *tangent space* T_pM is the set of equivalence classes $[\alpha]$, an n-dimensional vector space. Equivalently, T_pM is the space of *derivations* at p: $\mathbb{R}$-linear operators X on $C^\infty(M)$ satisfying the Leibniz rule $X(fg) = X(f)\, g(p) + f(p)\, X(g)$.

Observation 1.16. A tangent vector is the velocity of a curve through p. Without an ambient space one cannot write $\alpha'(0)$ as an arrow in $\mathbb{R}^3$; one reads it in a chart, $(\varphi \circ \alpha)'(0) \in \mathbb{R}^n$. For a surface in $\mathbb{R}^3$ this intrinsic definition coincides with the geometric tangent plane $d\mathbf{x}_q(\mathbb{R}^2)$ of Section 1.1. The thesis exploits exactly the local identification $T_pM \cong \mathbb{R}^2$ to fit local geometries on the flat domain.

1.3 First fundamental form

Definition 1.17 (First fundamental form)**.** The *first fundamental form* at p is the symmetric, positive-definite bilinear form $I_p : T_pS \times T_pS \to \mathbb{R}$ obtained by restricting the Euclidean inner product, $I_p(v, w) = \langle v, w\rangle$. In the basis $\{\mathbf{x}_u, \mathbf{x}_v\}$ its matrix is

$$(g_{ij}) = \begin{pmatrix} E & F \\ F & G \end{pmatrix}, \qquad E = \langle \mathbf{x}_u, \mathbf{x}_u\rangle, \quad F = \langle \mathbf{x}_u, \mathbf{x}_v\rangle, \quad G = \langle \mathbf{x}_v, \mathbf{x}_v\rangle. \tag{1.8}$$

The smooth assignment $p \mapsto I_p$ is a *Riemannian metric* g on S.

Remark 1.18. The coefficients E, F, G allow lengths, angles, and areas to be computed while staying on the surface; these are the *intrinsic* quantities. A length, for instance, is $\int \sqrt{E\,u'^2 + 2F\,u'v' + G\,v'^2}\,\mathrm{d}t$, and an area is $\iint \sqrt{EG - F^2}\,\mathrm{d}u\,\mathrm{d}v$. Everything the connection and the geodesics need is encoded in g.

Theorem 1.19 (Theorema Egregium)**.** The *Gaussian curvature* K of a surface (classically the product of its two principal curvatures) is *intrinsic*: it can be expressed through the coefficients E, F, G of the first fundamental form and their derivatives alone, and is therefore invariant under local isometries. Intrinsic, isometry-invariant quantities such as this are exactly the ones a shape-matching pipeline can rely on.

Observation 1.20. A flat sheet ($K \equiv 0$) can be rolled into a cylinder ($K \equiv 0$) without stretching, but cannot be wrapped onto a sphere ($K > 0$) without altering distances. This is why every map of the Earth distorts something, and why the flattening of a curved vertex star cannot be isometric.

1.4 Geodesics and the exponential map

The last tools we need translate between a curved surface and the flat tangent plane. They are built on *geodesics*, the curves that play on S the role of straight lines in $\mathbb{R}^n$, and they are the smooth model for the per-vertex constructions used later.

Definition 1.21 (Geodesic)**.** A curve γ on S is a *geodesic* if it has zero geodesic curvature, i.e. its acceleration is everywhere normal to the surface. In coordinates $\gamma(t) = \mathbf{x}(u^1(t), u^2(t))$ this is the second-order system

$$\ddot{u}^{\,k} + \Gamma^k_{ij}\,\dot{u}^{\,i}\,\dot{u}^{\,j} = 0, \qquad k = 1, 2, \tag{1.9}$$

where the $\Gamma^k_{ij} = \frac{1}{2}\,g^{km}(\partial_i g_{jm} + \partial_j g_{im} - \partial_m g_{ij})$ are the *Christoffel symbols* of the first fundamental form (Section 1.3), a repeated index being summed. By the existence and uniqueness theorem for ODEs, given $p = \gamma(0)$ and $v = \gamma'(0)$ there is a unique maximal geodesic, traversed at constant speed.

Observation 1.22 (Geodesics generalize straight lines)**.** A line in $\mathbb{R}^n$ has zero acceleration, $\ddot{\gamma} = 0$; a geodesic has zero *tangential* acceleration. It is the same equation, with the Γ^k_{ij} accounting for curvature: where they vanish, (1.9) reduces to $\ddot{u}^k = 0$, a straight line. On the sphere the geodesics are the great circles.

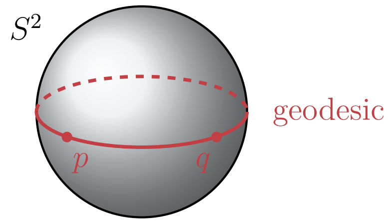

Figure 6. On the sphere the geodesics are the great circles, while in the plane they reduce to the straight lines.

The length of the shortest geodesic between two points defines the *geodesic distance* $d(p, q)$, the intrinsic notion of distance on S; it is the metric against which correspondences between shapes are scored in Chapter 5.

Definition 1.23 (Exponential and logarithmic maps)**.** For $v \in T_pS$ let γ_v be the unique geodesic with $\gamma_v(0) = p$ and $\gamma_v'(0) = v$. The *exponential map* is

$$\exp_p(v) = \gamma_v(1), \tag{1.10}$$

defined on a neighbourhood of the origin of T_pS. Its differential at the origin is the identity, so by the inverse function theorem $\exp_p$ is a diffeomorphism from a neighbourhood of 0 onto a normal neighbourhood of p. Its inverse is the *logarithmic map* $\log_p = \exp_p^{-1}$, and $\|\log_p(q)\| = d(p, q)$. Within the *injectivity radius* of p — bounded by the distance to the cut locus, where geodesics cease to be minimizing — both maps are well defined.

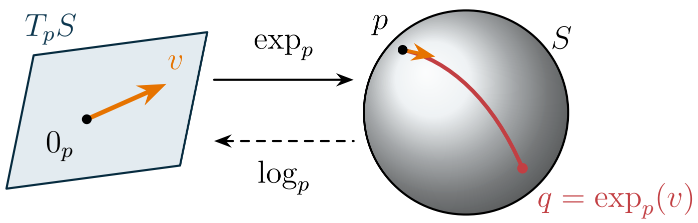

Figure 7. $\exp_p$ maps $v \in T_pS$ to the endpoint $q = \gamma_v(1)$ of the geodesic with initial velocity v; $\log_p$ maps q back to its tangent coordinate.

Remark 1.24 (The bridge to the discrete setting)**.** The pair exp / log translates between the tangent plane, where Euclidean tools apply, and the surface. This is the cornerstone of the constructions that follow: the smooth atlas at the centre of this thesis builds *discrete* exponential and logarithmic maps on a triangle mesh, flattening a vertex star onto the tangent plane while preserving the distance from the centre and the angular proportions. Curvature controls how much exp / log distort distances and angles away from p, and is the reason the flattening of a non-flat star cannot be isometric.

2. Spectral geometry

The previous chapter described a surface through the language of the smooth manifold: metric, tangent spaces, geodesics. That viewpoint is *local*. The constructions of this thesis instead rest on a *global*, analytic object attached to the metric — the Laplace–Beltrami operator — and on the orthonormal basis of functions it generates. This basis plays, on a curved surface, the role that the Fourier basis plays on the circle: it sorts functions by frequency and gives a compact, multi-scale way of representing both signals and maps between shapes. It is the basis on which the functional-map representation of Part II is built.

The presentation follows [4, 5, 6] and the classical spectral theory [7, 8].

2.1 Function spaces and the L^2 inner product

Let (M, g) be a compact, connected Riemannian surface without boundary (a closed surface, such as the human bodies of the FAUST dataset used in Part II). The Riemannian metric g of Section 1.3 fixes not only lengths and angles but also a canonical *area measure*: in local coordinates

$$d\mu = \sqrt{\det g}\, du^1\, du^2, \qquad \sqrt{\det g} = \sqrt{EG - F^2}. \tag{2.1}$$

Definition 2.1 (The space $L^2(M)$ and its inner product)**.** The space $L^2(M)$ consists of the (measurable) real functions $f : M \to \mathbb{R}$ with $\int_M f^2\, d\mu < \infty$. It is a Hilbert space for the L^2 *inner product*

$$\langle f, h\rangle_{L^2} = \int_M f(x)\, h(x)\, d\mu(x), \qquad \|f\|_{L^2}^2 = \langle f, f\rangle_{L^2}. \tag{2.2}$$

Smooth functions $C^\infty(M)$ are dense in it.

Remark 2.2. Equation (2.2) is the single most important quantity of the whole pipeline. Every projection onto the basis introduced below, every notion of "how much" two signals agree, and the discrete mass matrix M of Chapter 3 are concrete incarnations of this integral. Because $d\mu$ is built from the metric, the inner product is *intrinsic*: it does not see how M sits in $\mathbb{R}^3$.

Definition 2.3 (Gradient and divergence)**.** The *gradient* $\nabla f \in \mathfrak{X}(M)$ of $f \in C^\infty(M)$ is the unique vector field representing the differential through the metric, $\langle \nabla f, X\rangle_g = \mathrm{d}f(X)$ for every field X; in coordinates $(\nabla f)^i = g^{ij}\partial_j f$. The *divergence* $\operatorname{div} X$ of a field X is the function $\operatorname{div} X = \frac{1}{\sqrt{\det g}}\,\partial_i\Big(\sqrt{\det g}\,X^i\Big)$, the infinitesimal rate at which the flow of X expands area.

2.2 The Laplace–Beltrami operator

Definition 2.4 (Laplace–Beltrami operator)**.** The *Laplace–Beltrami operator* on (M, g) is the second-order differential operator

$$\Delta f = -\operatorname{div}(\nabla f) = -\frac{1}{\sqrt{\det g}}\,\partial_i\Big(\sqrt{\det g}\;g^{ij}\,\partial_j f\Big). \tag{2.3}$$

It generalizes to a curved surface the ordinary Laplacian $-\Big(\partial_{xx} + \partial_{yy}\Big)$ of the plane, to which it reduces when g is the identity.

Observation 2.5 (Sign convention)**.** We adopt the *geometer's sign*, with the minus sign in (2.3), so that Δ is positive semidefinite and its eigenvalues are non-negative. Analysts often write $\Delta = \operatorname{div}\nabla$ with the opposite sign; the two differ only by an overall sign and we shall always order eigenvalues so that $0 = \lambda_0 \le \lambda_1 \le \cdots$.

Proposition 2.6 (Self-adjointness and Green's identity)**.** On a closed surface Δ is self-adjoint and positive semidefinite with respect to (2.2): for all $f, h \in C^\infty(M)$,

$$\langle \Delta f, h\rangle_{L^2} = \langle \nabla f, \nabla h\rangle_{L^2} = \langle f, \Delta h\rangle_{L^2}, \qquad \langle \Delta f, f\rangle_{L^2} = \|\nabla f\|^2_{L^2} \ge 0. \tag{2.4}$$

The kernel of Δ consists of the constant functions.

Proof. The proof is based on the divergence theorem, which on a closed surface has no boundary term: for every smooth vector field X on M,

$$\int_M \operatorname{div} X \,\mathrm{d}\mu \;=\; 0. \tag{2.5}$$

Apply it to the field $X = h\,\nabla f$, whose divergence obeys the Leibniz rule

$$\operatorname{div}(h\,\nabla f) \;=\; h\,\operatorname{div}(\nabla f) + \langle \nabla h, \nabla f\rangle_g \;=\; -\,h\,\Delta f + \langle \nabla f, \nabla h\rangle_g,$$

the sign coming from the convention $\Delta = -\operatorname{div}\nabla$. Integrating and using (2.5) gives

$$0 \;=\; \int_M \Big(-\,h\,\Delta f + \langle \nabla f, \nabla h\rangle_g\Big)\mathrm{d}\mu, \qquad \text{that is} \qquad \langle \Delta f, h\rangle_{L^2} = \langle \nabla f, \nabla h\rangle_{L^2}.$$

The middle expression is symmetric in f and h, so repeating the argument with $X = f\,\nabla h$ yields $\langle \Delta f, h\rangle_{L^2} = \langle f, \Delta h\rangle_{L^2}$ (the operator is self-adjoint).

Taking $h = f$ gives $\langle \Delta f, f \rangle_{L^2} = \|\nabla f\|_{L^2}^2 \geq 0$, whence positive semidefiniteness. For the kernel, suppose $\Delta f = 0$. Then $\|\nabla f\|_{L^2}^2 = \langle \Delta f, f \rangle_{L^2} = 0$, so $\nabla f \equiv 0$ and f is locally constant; since M is connected, f is constant. The converse is immediate. Note that this last step is where connectedness is used: on a surface with several components the kernel would have one dimension per component. □

Remark 2.7 (Intrinsic and isometry-invariant)**.** Formula (2.3) involves only the metric coefficients g_{ij} and their derivatives. Hence Δ is *intrinsic* and, crucially, *commutes with isometries*: if $T : (M, g) \to (N, h)$ is an isometry and $f \in C^\infty(N)$, then $\Delta_M(f \circ T) = (\Delta_N f) \circ T$. This single fact is the reason the operator is so well suited to non-rigid shape matching: poses of the same subject are near-isometric, so their Laplacians — and the bases below — are nearly intertwined by the unknown map. We return to this in Part II as the *Laplacian-commutativity* constraint.

2.3 Spectrum and the eigenbasis

Theorem 2.8 (Spectral theorem on a closed surface)**.** Let M be a closed Riemannian surface. The Laplace–Beltrami operator Δ has a purely discrete spectrum

$$0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \cdots \longrightarrow +\infty, \tag{2.6}$$

and the associated eigenfunctions $\{\phi_k\}_{k\geq 0}$, defined by $\Delta \phi_k = \lambda_k \phi_k$, can be chosen to form an *orthonormal basis* of $L^2(M)$:

$$\langle \phi_j, \phi_k \rangle_{L^2} = \delta_{jk}, \qquad f = \sum_{k\geq 0} \hat{f}_k \, \phi_k, \quad \hat{f}_k = \langle f, \phi_k \rangle_{L^2}. \tag{2.7}$$

Observation 2.9. Expansion (2.7) is the exact analogue of a Fourier series, with $\{\phi_k\}$ replacing $\{\sin, \cos\}$ and $\{\hat{f}_k\}$ the generalized Fourier coefficients. The eigenvalue λ_k measures *frequency*: small λ_k correspond to slowly varying, smooth eigenfunctions, large λ_k to rapidly oscillating ones (Figure 8). On the unit circle the construction returns precisely the classical trigonometric basis with $\lambda_k = k^2$.

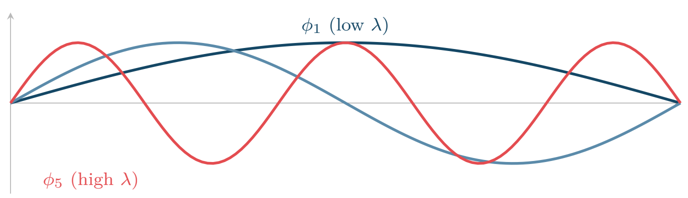

Figure 8. Eigenfunctions of the Laplacian on the interval $[0, 1]$, here $\phi_k(x) = \sin(k\pi x)$ with $\lambda_k = (k\pi)^2$. Increasing the eigenvalue increases the number of oscillations: the basis is naturally ordered from coarse to fine.

Definition 2.10 (Spectral truncation)**.** Fixing k, the *truncated basis* $\Phi = (\phi_0, \dots, \phi_{k-1})$ spans the k-dimensional space of the lowest-frequency functions. The truncation $f \approx \sum_{i<k} \hat{f}_i \, \phi_i$ is the best k-term approximation of f in L^2, and discarding high frequencies acts as a low-pass filter.

Remark 2.11 (Compactness and stability)**.** Two properties make this basis the natural choice for shape analysis [5, 4]. *Compactness*: because most signals of interest are smooth, a few dozen coefficients already reconstruct them well, so the truncation order k stays small (we use $k = 30$ in the experiments). *Stability*: although an *individual* eigenfunction is unstable — under perturbations of the shape it may flip sign or swap order with a neighbour when eigenvalues are close — the *subspace* spanned by the first k eigenfunctions is stable as long as λ_k and λ_{k+1} are well separated [8]. Functional maps (Part II) are designed to act between these stable subspaces rather than on the fragile eigenfunctions themselves.

2.4 Heat diffusion and spectral descriptors

The spectrum does more than supply a basis: it linearizes heat flow, and the heat flow in turn yields the point descriptors that drive the matching pipeline.

Definition 2.12 (Heat equation and heat kernel)**.** The diffusion of heat on M is governed by $\partial_t u + \Delta u = 0$. Its solution is given by the *heat kernel* $k_t(x, y)$, the amount of heat at y after time t produced by a unit source placed at x, which admits the spectral expansion

$$k_t(x, y) = \sum_{k \geq 0} e^{-\lambda_k t} \, \phi_k(x) \, \phi_k(y). \tag{2.8}$$

Observation 2.13. The factor $e^{-\lambda_k t}$ damps high frequencies exponentially fast, the more so as t grows. Small t keeps fine, local detail (governed by curvature near x); large t retains only the first eigenfunctions and thus the global shape. Diffusion time is therefore a continuous *scale* parameter.

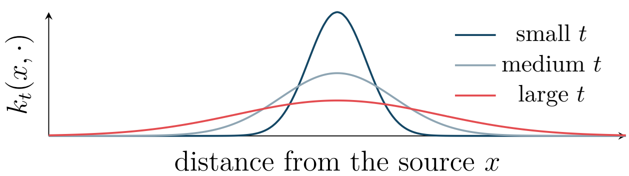

Figure 9. Heat kernels $k_t(x, \cdot)$ from a unit source at x for increasing diffusion times (shown on a line). Small t is sharp and local, capturing fine detail and curvature; large t spreads out and retains low-frequency, global structure. Diffusion time is the continuous scale sampled by the Heat Kernel Signature (2.9).

Definition 2.14 (Heat Kernel Signature)**.** The *Heat Kernel Signature* (HKS) [6] of a point x is the restriction of the heat kernel to the diagonal, sampled at a range of times:

$$\mathrm{hks}(x) = \Big(k_t(x,x)\Big)_t, \qquad k_t(x,x) = \sum_{k\geq 0} e^{-\lambda_k t}\,\phi_k(x)^2. \tag{2.9}$$

It associates with each point a vector that fingerprints its geometry across scales.

Remark 2.15 (Why HKS is the right descriptor here)**.** The signature (2.9) is built entirely from $\{\lambda_k, \phi_k\}$, hence it is *intrinsic* and *isometry-invariant*: corresponding points on two near-isometric poses receive almost the same signature. It is also informative and multi-scale, and provably stable [6]. An alternative spectral descriptor, the Wave Kernel Signature [10], replaces the heat kernel's low-pass smoothing with a band-pass response that better localises features; we keep HKS here because its multi-scale, low-pass character is the better match for the near-isometric matching of Part II. These are exactly the properties a descriptor must have to constrain a functional map by descriptor preservation (Part II). When the signature is conditioned on fixed source points one obtains the *landmark* HKS, the heat diffused from a chosen set of corresponding points [9]. The experiments of Part II use HKS and landmark HKS exactly as defined here, computed on a mesh by the discrete machinery of the next chapter.

3. Discrete geometry processing

The smooth theory of Chapters 1 and 2 lives on an ideal surface; computation lives on a *triangle mesh*, a piecewise-linear approximation of it. This chapter translates the continuous tools — the L^2 inner product, the Laplace–Beltrami operator, its eigenbasis, and the descriptors — into the finite-dimensional matrices that the code of Part II actually manipulates. It also introduces the surface sampling and the linear (barycentric) interpolation that serve as the *baseline* against which the manifold-based interpolation is measured, and the geodesic error used to score a correspondence.

The discretization follows [11, 12] and the evaluation protocol [13].

3.1 Triangle meshes and piecewise-linear functions

Definition 3.1 (Triangle mesh)**.** A *triangle mesh* is a pair $\mathcal{M} = (V, F)$ with vertices $V = \{v_1, \ldots, v_n\} \subset \mathbb{R}^3$ and triangular faces $F \subseteq V^3$, together with the piecewise-flat surface obtained by filling each face with a Euclidean triangle. The *one-ring* (or *star*) of a vertex v_i is the set of faces incident to it; two vertices are *neighbours*, $i \sim j$, if they share an edge.

Definition 3.2 (Piecewise-linear functions and barycentric coordinates)**.** A function is represented by its values $f = (f_1, \ldots, f_n)^\top \in \mathbb{R}^n$ at the vertices and extended *linearly* inside each face. A point p in a face (v_a, v_b, v_c) is written through its *barycentric coordinates* $(\beta_a, \beta_b, \beta_c)$, with $\beta_a, \beta_b, \beta_c \geq 0$ and $\beta_a + \beta_b + \beta_c = 1$, as $p = \beta_a v_a + \beta_b v_b + \beta_c v_c$; the interpolated value is

$$f(p) = \beta_a f_a + \beta_b f_b + \beta_c f_c. \tag{3.1}$$

Equivalently, $\{f_i\}$ are the coefficients of f in the *hat basis* of the piecewise-linear functions $\{\psi_i\}$, where $\psi_i(v_j) = \delta_{ij}$.

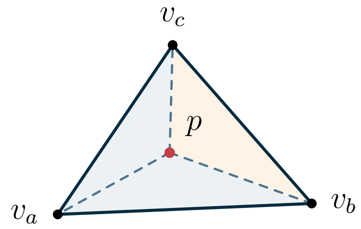

$$p = \beta_a v_a + \beta_b v_b + \beta_c v_c, \qquad \beta_a + \beta_b + \beta_c = 1$$

Figure 10. A point p inside a triangle and its barycentric coordinates.

3.2 Discrete Laplacian and mass matrix

The continuous identities $\langle f, h\rangle_{L^2} = \int_M fh\,\mathrm{d}\mu$ and $\langle \Delta f, h\rangle_{L^2} = \int_M \nabla f \cdot \nabla h\,\mathrm{d}\mu$ become, on the hat basis, two sparse $n \times n$ matrices.

Definition 3.3 (Mass and stiffness matrices)**.** The *mass matrix* M discretizes the L^2 inner product, $\langle f, h\rangle_{L^2} \approx f^\top M\,h$. In the common *lumped* form it is diagonal, $M_{ii} = A_i$, with A_i the area attached to vertex v_i (a third of the area of its incident triangles, or the local Voronoi area). The *stiffness matrix* W discretizes the Dirichlet energy through the *cotangent weights* [12, 11]

$$W_{ij} = -\tfrac{1}{2}\Big(\cot\alpha_{ij} + \cot\beta_{ij}\Big)\ (i \sim j), \qquad W_{ii} = -\sum_{j\sim i} W_{ij}, \tag{3.2}$$

where α_{ij}, β_{ij} are the two angles opposite the edge (i, j) (Figure 11). We reserve Δ for the smooth operator and denote its matrix discretisation by L. The discrete Laplace–Beltrami operator is then

$$L = M^{-1}W, \qquad f^\top W f = \tfrac{1}{2}\sum_{i\sim j}\Big(\cot\alpha_{ij} + \cot\beta_{ij}\Big)(f_i - f_j)^2 \geq 0. \tag{3.3}$$

The cotangents in (3.2) look arbitrary, and it is worth seeing that they are not: they are forced, with no modelling freedom, by requiring that the discrete Dirichlet energy be the exact energy of the piecewise-linear interpolant. The following classical computation [12, 11] derives them.

Proposition 3.4 (The cotangent formula)**.** Let f be piecewise linear on the mesh, with values f_i at the vertices. Then its Dirichlet energy is

$$E(f) \;=\; \frac{1}{2}\int_{\mathcal{M}} \|\nabla f\|^2 \;=\; \frac{1}{4}\sum_{i\sim j}\Big(\cot\alpha_{ij} + \cot\beta_{ij}\Big)\,(f_i - f_j)^2,$$

so that $E(f) = \frac{1}{2}\,f^\top W f$ with W as in (3.2).

Proof. Since f is linear on each triangle, ∇f is constant there and the integral splits over triangles: it suffices to compute the contribution of a single triangle $t = t_{ijk}$ of area A_t.
On t the gradient of the hat function ϕ_i is the vector orthogonal to the opposite edge $e_i = v_k - v_j$, pointing towards v_i, with magnitude the reciprocal of the corresponding height. Writing $e_i^\perp$ for the rotation of e_i by $\pi/2$ in the plane of the triangle,

$$\nabla\phi_i \;=\; \frac{e_i^\perp}{2A_t}, \qquad \text{hence} \qquad \nabla f \;=\; \sum_{m\in\{i,j,k\}} f_m\,\nabla\phi_m \;=\; \frac{1}{2A_t}\sum_m f_m\,e_m^\perp.$$

Because $e_i + e_j + e_k = 0$, the same holds for the rotated edges, so the coefficients may be shifted by a constant without changing ∇f — the gradient sees only the *differences* of the f_m, as it must. Therefore

$$2\,E_t(f) \;=\; A_t\,\|\nabla f\|^2 \;=\; \frac{1}{4A_t}\Big\|\sum_m f_m\,e_m^\perp\Big\|^2 \;=\; \frac{1}{4A_t}\sum_{m,n} f_m f_n\,\langle e_m, e_n\rangle,$$

the rotation being an isometry. Using $\sum_m e_m = 0$ to eliminate the diagonal terms, $\langle e_m, e_m\rangle = -\sum_{n\neq m}\langle e_m, e_n\rangle$, the double sum collapses to a sum over the three pairs:

$$2\,E_t(f) \;=\; -\frac{1}{4A_t}\sum_{m<n}\langle e_m, e_n\rangle\,(f_m - f_n)^2.$$

It remains to identify the coefficient. If θ_k is the angle of t at the vertex v_k, the two edges meeting there are e_i and e_j up to sign, and

$$\langle e_i, e_j\rangle = -\|e_i\|\|e_j\|\cos\theta_k, \qquad 2A_t = \|e_i\|\|e_j\|\sin\theta_k,$$

so that $-\langle e_i, e_j\rangle/(2A_t) = \cot\theta_k$. Substituting,

$$E_t(f) \;=\; \frac{1}{4}\sum_{m<n}\cot\theta_{mn}\,(f_m - f_n)^2,$$

where θ_{mn} is the angle of t opposite the edge (m,n). Summing over all triangles, each interior edge (i,j) is shared by exactly two of them and collects the two opposite angles α_{ij} and β_{ij}, which gives the stated formula. $\square$

Observation 3.5 (Why the weights can be negative)**.** Nothing in the derivation forces $\cot\alpha_{ij} + \cot\beta_{ij}$ to be positive: it is negative exactly when the two angles opposite an edge sum to more than π, i.e. on badly shaped triangles. The energy $f^\top W f$ then remains non-negative — it is an exact energy, and W is positive semidefinite by construction — but individual weights are not, and the discrete operator loses the maximum principle. This is one concrete way in which a poor tessellation degrades a discrete operator, and it is worth keeping in mind alongside the failure of the angle defect on the irregular meshes of Chapter 7.

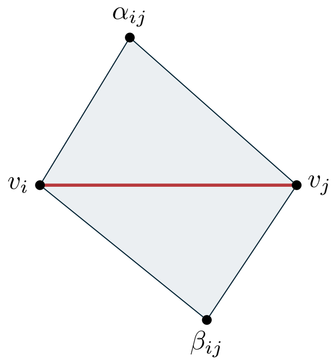

Figure 11. The interior edge (v_i, v_j) (red) is shared by two triangles; the cotangent weight (3.2) is built from the two angles α_{ij}, β_{ij} opposite it. These weights make the discrete Laplacian a faithful, mesh-aware analogue of Δ.

Observation 3.6 (The discrete eigenproblem)**.** Eigenfunctions are computed by solving the *generalized* eigenproblem

$$W\phi_k = \lambda_k \, M \, \phi_k, \qquad \Phi^\top M \, \Phi = I, \tag{3.4}$$

which yields M-orthonormal discrete eigenvectors $\Phi = [\phi_0, \ldots, \phi_{k-1}] \in \mathbb{R}^{n \times k}$ and eigenvalues $\Lambda = \mathrm{diag}(\lambda_0, \ldots, \lambda_{k-1})$. The M-orthonormality in (3.4) is the discrete counterpart of $\langle \phi_j, \phi_k \rangle_{L^2} = \delta_{jk}$.

Remark 3.7 (Spectral projection)**.** With Φ and M in hand, the generalized Fourier coefficients (2.7) of a vertex signal $f \in \mathbb{R}^n$ are the matrix–vector product

$$\hat{f} = \Phi^\top M f \in \mathbb{R}^k. \tag{3.5}$$

This *projection* $\Phi^\top M f$ is the operation that the functional-map pipeline of Part II applies to every descriptor; the entire difference between the baseline and the manifold-based experiments amounts to changing how the quantities entering (3.5) are evaluated.

3.3 Surface sampling and quadrature

The projection (3.5) integrates over the mesh through the vertex mass matrix. An alternative, used throughout the experiments, replaces that integral by a *quadrature* over points scattered on the surface — which is what lets a continuous, non-vertex representation enter the same pipeline.

Definition 3.8 (Area-weighted surface sampling)**.** Draw N points $\{x_1, \ldots, x_N\}$ on the mesh by selecting each face with probability proportional to its area and a uniform barycentric point within it. Each sample carries an *area weight* ω_i with $\sum_i \omega_i \approx \mathrm{Area}(\mathcal{M})$; write $\omega = (\omega_1, \ldots, \omega_N)^\top$. Any surface integral is then approximated by the Monte-Carlo quadrature

$$\int_M f \,\mathrm{d}\mu \;\approx\; \sum_{i=1}^{N} \omega_i\, f(x_i). \tag{3.6}$$

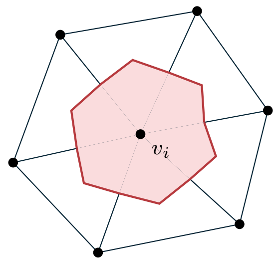

Figure 12. The area A_i of vertex v_i is its barycentric dual cell (red): one third of each incident triangle. It is the diagonal entry M_{ii} of the lumped mass matrix and, equivalently, the area weight ω_i of the quadrature (3.6).

Observation 3.9 (Two ways to project the same signal)**.** Equation (3.6) turns the spectral projection into

$$\hat{f} \;\approx\; \Phi(X)^\top \,\mathrm{diag}(\omega)\, f(X), \tag{3.7}$$

where $\Phi(X) \in \mathbb{R}^{N\times k}$ collects the eigenfunctions at the sample points and $f(X)$ the signal there. Comparing the vertex projection (3.5) with the sampled one (3.7) is the methodological heart of the experiments: the diagonal weight matrix $\mathrm{diag}(\omega)$ plays, at the sample points, the role of the mass matrix M.

Remark 3.10 (Linear interpolation as the baseline)**.** To evaluate $\Phi(X)$ and $f(X)$ one must read vertex quantities at off-vertex points. The elementary choice is the linear, barycentric interpolation (3.1) of Section 3.1. It is exact only for functions that are affine inside each triangle and therefore blurs the smooth structure of the true eigenfunctions and descriptors. Part II replaces precisely this step by a manifold-based interpolation that respects that structure, and asks whether the same accuracy can then be reached with *fewer* samples N.

3.4 Spectral descriptors on a mesh

Definition 3.11 (Discrete HKS and landmark HKS)**.** From the discrete spectrum (Φ, Λ) the Heat Kernel Signature (2.9) of vertex v_i over a set of diffusion times t is

$$\mathrm{hks}_t(v_i) = \sum_k e^{-\lambda_k t}\, \Phi_{ik}^2, \tag{3.8}$$

giving a matrix $\mathrm{HKS} \in \mathbb{R}^{n_t \times n}$. The *landmark* HKS repeats the construction for heat sources placed at a chosen set of landmark vertices, stacking one block per landmark.

Remark 3.12 (The descriptor pipeline)**.** In practice HKS and landmark HKS are concatenated, subsampled in the time domain, and L^2-normalized, producing a stack of descriptor channels per shape. These are the functions f that, once projected by (3.5) or (3.7), constrain the functional map. Each channel is treated independently, which is exactly what makes descriptor preservation a *linear* constraint in Part II.

3.5 From maps to correspondences and their error

The output to be evaluated is ultimately a point-to-point correspondence between two meshes; its quality is measured intrinsically.

Definition 3.13 (Geodesic error)**.** Let $T : \mathcal{M} \to \mathcal{N}$ be a computed correspondence and $T^\star$ the ground-truth one. For a vertex v the *geodesic error* is the intrinsic distance on $\mathcal{N}$ between the two images, normalized by the surface scale,

$$\varepsilon_{\mathrm{geo}}(v) = d_{\mathcal{N}}\Big(T(v),\, T^\star(v)\Big) \big/ \sqrt{\mathrm{Area}(\mathcal{N})}. \tag{3.9}$$

A perfect map has $\varepsilon_{\mathrm{geo}} \equiv 0$; the average of $\varepsilon_{\mathrm{geo}}$ over all vertices is the standard scalar score [13].

Remark 3.14 (Datasets and protocol)**.** The experiments use the FAUST dataset of registered human-body scans [14], which supplies dense ground-truth correspondences between poses, together with the evaluation protocol of [13]: convert the functional map to a point-to-point map, then report the average geodesic error (3.9). The geodesic distance itself is computed on the mesh by a discrete distance solver. With the mass matrix, the eigenbasis, the descriptors, the sampling, and this error in place, every ingredient that the functional-map experiments of Part II need is now expressed as a concrete, computable object.

$$\begin{tikzcd}
L^2(\mathcal{M}) \arrow[d, "\Phi_{\mathcal{M}}^\top M_{\mathcal{M}}"'] & L^2(\mathcal{N}) \arrow[l, "T_F"'] \arrow[d, "\Phi_{\mathcal{N}}^\top M_{\mathcal{N}}"] \\
\mathbb{R}^k & \mathbb{R}^k \arrow[l, "\mathbf{C}"]
\end{tikzcd}$$

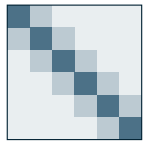

$\mathbf{C}$

Figure 13. A point-to-point map between two shapes induces a linear map T_F between their function spaces; written in the truncated Laplacian eigenbases, through the projection $\Phi^\top M$ of (3.5), it collapses to a small matrix, the *functional map* $\mathbf{C}$. For a near-isometry $\mathbf{C}$ is close to diagonal (right). Part II estimates $\mathbf{C}$ from descriptor constraints, converts it back to a correspondence, and scores it with the geodesic error (3.9).

4. DALa

The atlas that the experiments of Part II analyse is produced by **DALa** [19], a domain-specific language and accompanying atlas construction developed by Claudio Mancinelli and collaborators. The manuscript describing the language is still in preparation at the time of writing. This chapter gives a brief, self-contained account of the specific pieces on which the experimental chapters rely. The implementation remains externally supplied, but the mathematical properties used in Chapter 8 are stated explicitly rather than treated as black-box assumptions.

DALa — a *Differentiable Atlas Language* — is a domain-specific language for expressing differential-geometric computations directly on the smooth manifold implicitly defined by an atlas. A program written in it reads almost like the corresponding textbook formula: the surface normal, for instance, is written as the normalised cross product of the two parametric derivatives of the surface, and the language evaluates that expression on the atlas without the programmer ever manipulating transition functions or coding derivatives by hand. The engine underneath is *automatic differentiation* in forward mode: because the elementary chart and blending operations are differentiable and the differentiation operators can be nested, DALa propagates derivatives of arbitrary order through the whole surface-evaluation pipeline. This makes it possible to evaluate *higher-order* differential quantities — among them the Gaussian curvature this thesis studies — by automatic differentiation at any point of a triangle, without finite differences on a mesh.

The second ingredient DALa introduces, and the one that matters for us, is a new *atlas construction* tailored to that differentiable model. It follows the vertex-centric philosophy of Djuren et al. [18] — a local polynomial chart per vertex, blended into a global parametrisation — but extends it from G^1 to prescribed C^k regularity; a modified choice of weights also supports C^∞. The charts, transition maps and blending are designed so that the automatic-differentiation engine can generate consistent derivatives of the prescribed order. The precise geometric qualification matters: the construction first gives a C^k map on the abstract manifold associated with the atlas; it is an immersed parametric pseudo-manifold under the rank assumptions of Proposition 4.1, and its image is an embedded surface only when the global map is also injective.

4.1 The atlas construction, in brief

The construction is guided by four requirements:

- prescribed C^k regularity up to $k = \infty$;
- chart supports as compact as possible;
- a *geometry-aware* chart parametrisation;
- a well-conditioned evaluation everywhere.

It rests on two deliberately separated stages.

Local charts. Each vertex, edge and triangle of the mesh carries a planar chart domain. The vertex chart is obtained by *conformal flattening* of the vertex star: the incident triangles are flattened around the origin and their angles rescaled so that the total angle is 2π, while radial distances are preserved. A local polynomial — the *vertex polynomial*, of a chosen degree — is then fitted over this domain to reproduce the geometry in a neighbourhood of the vertex. This is the "guide-surface" idea of [18], promoted here to arbitrary degree and continuity.

Blending and the decay parameter. On a triangle chart U_T, every active vertex patch is first pulled back to the same coordinates through the transition map $\varphi_{T,i}$. We denote the global reconstructed map by S and its expression on U_T by $S_T := S|_{U_T}$. The local parametrisation is therefore

$$S_T(u) = \sum_i w_i(u)\, S_i(\varphi_{T,i}(u))\,, \qquad \sum_i w_i(u) = 1, \tag{4.1}$$

where the sum is over the vertices of T. This common-coordinate form is the object dissected in Chapter 8. DALa's contribution here is the shape of the weights w_i.

In [18] the vertex weight functions do not vanish at the boundary of the chart domains, which lets the *unbounded derivatives of the conformal flattening* near non-flat vertices leak into the blend and cause numerical instabilities. DALa introduces a parametric blending region governed by a *decay parameter* ε: the weights are forced to decay smoothly to zero *before* reaching the chart boundary, so that the local polynomial alone is responsible for the geometry near each vertex and the partition of unity is activated only in the overlap region away from the flattening singularity. This localisation keeps the parametrisation regular, shrinks the active blending region (a sparser computational graph for the differentiation engine), and is precisely the mechanism whose second-order side effects Chapter 8 analyses.

More precisely, let $\bar{w}_i$ denote the non-negative, edge-compatible weights before compactification. For $0 < \varepsilon < 1/3$, DALa defines

$$\eta_k(t) = \begin{cases} t^{k+1}, & t > 0, \\ 0, & t \leq 0, \end{cases} \qquad g_{\varepsilon,k}(s) = s\,\eta_k(s - \varepsilon),$$

and normalises the compactified weights,

$$w_i(u) = \frac{g_{\varepsilon,k}(\bar{w}_i(u))}{\sum_j g_{\varepsilon,k}(\bar{w}_j(u))}. \tag{4.2}$$

The C^∞ variant replaces η_k with $\eta_\infty(t) = \exp(-1/t)$ for $t > 0$ and zero otherwise. Thus ε is a threshold in the parametric domain, not itself a physical bandwidth. Increasing it enlarges the single-chart plateaus and reduces the region in which several patches are blended.

The vertex plateau. For a vertex i of T, define

$$N^{\varepsilon}_{T,i} = \{u \in U_T : \bar{w}_j(u) \leq \varepsilon \text{ for every } j \neq i\}\,.$$

On this open neighbourhood the compactified weights satisfy $w_i = 1$ and $w_j = 0$ for $j \neq i$, with all their derivatives equal to zero. Consequently, Equation (4.1) reduces exactly to

$$S_T(u) = S_i(\varphi_{T,i}(u))\,, \qquad u \in N^{\varepsilon}_{T,i}. \tag{4.3}$$

This is a property of the atlas construction, not merely an optimisation in the curvature evaluator. Its role in the comparison with local polynomial fitting is discussed in Section 8.1.

Proposition 4.1 (Regularity of the DALa parametrisation)**.** Assume that every local shape function S_i is C^k. The weights above define compatible local parametrisations and hence a globally defined C^k map on the abstract manifold associated with the atlas. If the local patches have rank two on the vertex plateaus and the blended maps have rank two elsewhere, the map is an immersion and defines a C^k parametric pseudo-manifold. If it is also globally injective, its image is an embedded C^k surface.

The proposition guarantees compatibility and regularity; it makes no statement about how accurately the image approximates the unknown smooth surface, nor about convergence under refinement. Those are separate approximation questions studied empirically in Part II and discussed through a conjectural comparison in Chapter 8.

Fitting the charts. Collecting all polynomial coefficients in a vector C, DALa assembles stretch and bending matrices Q_s and Q_b by numerical quadrature. In the hard mode it solves the constrained quadratic problem

$$\underset{C}{\text{minimise}}\; C^\top(\lambda_s Q_s + \lambda_b Q_b)C \qquad \text{subject to} \qquad AC = b, \tag{4.4}$$

where $AC = b$ imposes interpolation of the mesh vertices. In the weak mode this constraint is replaced by a weighted data term,

$$\underset{C}{\text{minimise}}\; C^\top(\lambda_s Q_s + \lambda_b Q_b)C + \lambda_i \|AC - b\|_W^2. \tag{4.5}$$

A surface attractor analogously penalises deviation from a target sampled over the whole triangle. These formulations separate three possible sources of approximation error: the finite polynomial space, the variational bias introduced by the energy weights, and the quadrature used to assemble the matrices.

4.2 What the experiments use

Only a slice of DALa is exercised in Part II, and it is worth naming it explicitly.

The atlas as a curvature estimator. The experiments query, at a surface point p, the reconstructed position $S(p)$ and the Gaussian curvature $K_S(p)$ read from the first and second fundamental forms of the charts. Here the subscript in K_S only records that the same Gaussian curvature K is evaluated on the reconstructed surface S. It is DALa's automatic differentiation that supplies those fundamental forms directly; nothing is finite-differenced.

The fitting modes. The *strong*, *weak* and *surface-attractor* estimators of Chapter 7 are exactly the hard-interpolation, pointwise-attractor and surface-attractor options of the fitting problem above, with the PN-triangle or the analytic torus as the surface-attractor target (Experiment A4).

The blending weights. The gate experiment of Chapter 8 compares the two ways of applying the ε-governed gate — to the blending weight, which is what provides the proved C^k compatibility, or to the barycentric coordinate, which is softer but carries no such guarantee. The whole discussion of the localised curvature artefact is a study of the derivatives of these weights, and of the regularisation that damps their effect.

Degree and continuity. The reconstructions use degree-three vertex polynomials and implementation continuity parameter $k_{\text{cont}} = 2$ throughout, except in the degrees-of-freedom experiment, which raises the degree to five to probe the fit's dependence on the number of chart coefficients.

The reconstruction quadrature. The fixed quadrature that DALa uses to assemble the fitting energy is the object of the quadrature-ladder experiment, which shows that it can be a limiting factor at the tested fine resolution.

The remaining parts of DALa — its syntax, differentiation machinery, and applications beyond curvature — lie outside the scope of this thesis. What we borrow is the differentiable C^k atlas and the handful of parameters that shape it, and what we contribute is an account of how faithfully that atlas, fitted from positions alone, reproduces a second-order differential quantity.

5. Functional Maps

The preliminaries gave, for each shape, an intrinsic Laplacian eigenbasis (Chapter 2) and a set of intrinsic descriptors built from it (the HKS of Section 2.4). This chapter assembles them into the *functional map* of Ovsjanikov et al. [3], a compact linear encoding of a correspondence between two shapes. We follow their formulation, specialised to the intrinsic, mesh-based setting of Part I. This construction will support a secondary downstream experiment in Chapter 7.

5.1 From point maps to functional maps

Let $\mathcal{M}$ and $\mathcal{N}$ be two shapes and $T : \mathcal{M} \to \mathcal{N}$ a point-to-point correspondence between them. Rather than tracking where each point is sent, one tracks how *functions* are transported.

Definition 5.1 (Induced functional map)**.** A point map $T : \mathcal{M} \to \mathcal{N}$ induces a *functional map* $T_F : L^2(\mathcal{N}) \to L^2(\mathcal{M})$ by pull-back,

$$T_F(g) = g \circ T, \qquad g \in L^2(\mathcal{N}). \tag{5.1}$$

Observation 5.2 (Why this helps)**.** Whatever the (highly non-linear) map T is, the induced T_F is always *linear*: $T_F(\alpha g_1 + \beta g_2) = \alpha\, T_F(g_1) + \beta\, T_F(g_2)$. The search for a correspondence is thus moved from the combinatorial space of point maps to the linear space of operators between function spaces, where the tools of linear algebra (Appendix B) apply directly.

5.2 Matrix representation in a reduced basis

Linear operators are represented by matrices once a basis is fixed. The natural choice is the truncated Laplace–Beltrami eigenbasis: let $\Phi_{\mathcal{M}} = (\phi_1^{\mathcal{M}}, \ldots, \phi_k^{\mathcal{M}})$ and $\Phi_{\mathcal{N}} = (\phi_1^{\mathcal{N}}, \ldots, \phi_k^{\mathcal{N}})$ be the first k eigenfunctions of $\mathcal{M}$ and $\mathcal{N}$.

Definition 5.3 (Functional map matrix)**.** Expanding any $g \in L^2(\mathcal{N})$ as $g = \sum_j b_j\, \phi_j^{\mathcal{N}}$ and its image $T_F(g) = \sum_i a_i\, \phi_i^{\mathcal{M}}$, linearity gives $a = \mathbf{C}\, b$ with

$$\mathbf{C}_{ij} = \left\langle\, T_F(\phi_j^{\mathcal{N}}),\ \phi_i^{\mathcal{M}} \,\right\rangle_{L^2(\mathcal{M})}, \qquad \mathbf{C} \in \mathbb{R}^{k\times k}. \tag{5.2}$$

The single small matrix $\mathbf{C}$ *is* the functional map: it acts on the k-dimensional spectral coefficients, not on the n mesh vertices.

Remark 5.4 (Compactness)**.** Reducing the basis to its first k entries (a few tens in practice, as in Section 3.4) makes $\mathbf{C}$ a $k \times k$ matrix *independent of the mesh resolution*. This is the property that lets the representation be estimated, regularised and stored cheaply, and is what makes a fair comparison across different samplings possible at all.

Proposition 5.5 (Structure for (near-)isometries)**.** If T is an isometry, then T_F maps eigenfunctions to eigenfunctions of equal eigenvalue, so $\mathbf{C}$ is *orthonormal* and, when the spectrum is non-repeated, *diagonal* up to sign. For a near-isometry $\mathbf{C}$ stays close to a funnel-shaped, near-diagonal orthonormal matrix (Figure 13). Deviation from this structure is precisely what the estimation below penalises.

5.3 Estimating the functional map

In practice T is unknown and $\mathbf{C}$ is recovered from quantities that are preserved by the sought correspondence.

Definition 5.6 (Descriptor preservation)**.** Let $f_1, \dots, f_q$ be descriptors on $\mathcal{M}$ and $g_1, \dots, g_q$ the corresponding descriptors on $\mathcal{N}$. Stacking their spectral coefficients (computed by the projection $\Phi^\top M$ of (3.5)) as the columns of $\mathbf{A} = \Phi_{\mathcal{N}}^{+} G$ and $\mathbf{B} = \Phi_{\mathcal{M}}^{+} F$, a correspondence-preserving $\mathbf{C}$ should satisfy

$$\mathbf{C}\,\mathbf{A} \approx \mathbf{B}. \tag{5.3}$$

Observation 5.7 (Commutativity with the Laplacian)**.** An isometry commutes with the Laplace–Beltrami operator, which in the eigenbasis means $\mathbf{C}$ should nearly commute with the (diagonal) spectra $\Lambda_{\mathcal{M}}, \Lambda_{\mathcal{N}}$:

$$\mathbf{C}\,\Lambda_{\mathcal{N}} \approx \Lambda_{\mathcal{M}}\,\mathbf{C}. \tag{5.4}$$

This is a descriptor-free regulariser that pushes $\mathbf{C}$ towards the near-diagonal structure of Proposition 5.5.

Definition 5.8 (The functional-map energy)**.** The map is estimated as the least-squares minimiser

$$\mathbf{C}^\star = \arg\min_{\mathbf{C}} \underbrace{\|\mathbf{C}\,\mathbf{A} - \mathbf{B}\|_F^2}_{\text{descriptor preservation}} + \mu \underbrace{\|\mathbf{C}\,\Lambda_{\mathcal{N}} - \Lambda_{\mathcal{M}}\,\mathbf{C}\|_F^2}_{\text{commutativity}}, \tag{5.5}$$

a convex problem in the k^2 entries of $\mathbf{C}$. Without the regulariser it reduces to the linear least-squares problem $\mathbf{C} = \mathbf{B}\,\mathbf{A}^+$ of Appendix B.

Remark 5.9 (Where the functional-maps experiment enters)**.** Every term of (5.5) is built from two ingredients only: the eigenbasis Φ and the descriptor coefficients $\Phi^\top M f$. The functional-maps experiment of Chapter 7 will change *how those two ingredients are evaluated* while keeping (5.5) fixed.

5.4 From a f-map back to a correspondence

A functional map is only useful if it can be turned back into the point-to-point correspondence that is ultimately evaluated.

Definition 5.10 (Spectral embedding and recovery)**.** Each point $x \in \mathcal{M}$ is embedded as the row $\Phi_{\mathcal{M}}(x) \in \mathbb{R}^k$ of eigenfunction values at x; likewise for $\mathcal{N}$. A point map consistent with $\mathbf{C}$ is recovered by nearest neighbour in this embedding,

$$T(x) = \arg\min_{y \in \mathcal{N}} \left\| \mathbf{C}^\top \Phi_{\mathcal{M}}(x)^\top - \Phi_{\mathcal{N}}(y)^\top \right\|_2. \tag{5.6}$$

Remark 5.11 (Evaluation)**.** The recovered T is scored, against the ground-truth correspondences of the FAUST dataset [14], by the average geodesic error (3.9) under the protocol of [13]. With the estimation (5.5) and the recovery (5.6) in place, the functional-map pipeline is complete and ready to receive, interchangeably, each of the representations considered in the downstream experiment of Chapter 7.

PART II

Experiments

6. Motivation

The preliminaries built, one construction at a time, the smooth C^k atlas used throughout this thesis: a surface reconstruction that attaches a polynomial chart to every vertex of a triangle mesh and blends the charts into a globally C^k parametrisation. The method was developed by Prof. Claudio Mancinelli and collaborators and is described in the DALa manuscript [19]; here we focus on its differential behaviour while treating the externally supplied implementation as fixed. Two properties of that construction, taken together, raise the question studied in this thesis.

6.1 A tool built for positions

The atlas is fitted with *order-zero information only.* Its charts are constrained, softly or exactly, to reproduce the *positions* of the mesh vertices; nothing in the fit prescribes normals, curvature, or any higher-order differential quantity. In this respect it belongs to the same family as the osculating-jet fits of Cazals and Pouget [17] and the quadratic fitting of Xu [16]: a surface is recovered from sampled points, and whatever differential information it carries is a *by-product* of the fit, not an input to it. What sets the DALa atlas apart is its prescribed compatibility and the fact that its polynomial charts can be differentiated by automatic differentiation, without finite differences, at *any* point of a triangle and not only at its vertices. Classical mesh-based estimators of curvature — the angle-defect Gaussian curvature, the cotangent Laplacian, the discrete mean-curvature vector — are defined *at the vertices* and read the piecewise-linear mesh directly; they have no surface to differentiate away from the vertices, and no notion of a higher-order derivative beyond the combinatorial one. The atlas, having a smooth parametrisation underneath, does. Yet the construction was designed as a *positional* reconstruction: its fitting objective measures how faithfully it reproduces the surface, not how faithfully it reproduces the surface's derivatives.

The question is therefore natural and, we believe, new:

> *A smooth atlas reconstructs a surface accurately in position, using only positional data. Does it, as a consequence, reconstruct the surface's higher-order differential quantities accurately as well?*

The question becomes practically accessible because the atlas provides a recent, globally compatible reconstruction whose derivatives can be evaluated throughout each triangle, rather than only pointwise at mesh vertices.

Why curvature. We answer the question for the Gaussian curvature. The choice is not incidental. Curvature is the prototypical second-order differential quantity of a surface, and one of the most heavily used in geometry processing: it drives feature detection, remeshing, segmentation, smoothing and fairing, registration, and the discretisation of the geometric operators (Laplace–Beltrami, Giaquinta–Hildebrandt) on which a large part of the field rests [16, 11]. If a smooth atlas is to be more than a prettier interpolant — if its differentiability is to be an asset rather than a decoration — then the curvature it produces is the first thing that ought to be trustworthy.

It is also the place where a positional reconstruction is most likely to disappoint. Position is an order-zero quantity and the fit controls it directly; curvature is an order-two quantity and the fit never sees it. Between the two lies room for a reconstruction to be simultaneously faithful in shape and wrong in curvature. Whether that room is actually occupied is an empirical question, and Chapter 7 tests it directly.

Why the torus. To measure a curvature error one needs a surface whose curvature is known in closed form, so that the estimate can be compared against ground truth rather than against another estimate. This already restricts the testbed to the few surfaces with an analytic second fundamental form.

The obvious candidate is the sphere, and we deliberately avoid it. Its constant curvature and high symmetry make it an unusually favourable test case and can conceal directional or sign-dependent errors. Moreover, the atlas fit includes a bending regulariser, so testing only a sphere would make it difficult to separate geometric accuracy from the regularising preference of the objective.

The torus of revolution provides a less symmetric test while keeping the analytic ground truth. Its Gaussian curvature is known exactly and, crucially, it is *signed and varying*: positive on the outer equator, negative on the inner one, zero along the top and bottom circles. A single torus thus exercises the estimator across the whole range of curvature signs and magnitudes. Its inner region, where the curvature is negative, probes behaviour that a constant-curvature test cannot expose. This combination makes it a deliberately demanding analytic testbed.

6.2 Empirical claims and scope

On this testbed the thesis makes three empirical claims, developed in Chapter 7 and interpreted in Chapter 8:

1. **selected atlas reconstructions improve under the tested refinements, including on irregular tessellations.** On irregular meshes, where the classical angle-defect error does not decrease, the strong, weak, and analytically targeted reconstructions retain a decreasing error trend. This is the positive empirical result. It is qualitatively consistent with Xu's analysis of a *local quadratic* fit [16], but Xu's theorem neither covers DALa's degree-three global variational solve nor supplies an asymptotic rate for it;
2. **under-resolved numerical integration can mask a decreasing error trend.** An apparent error floor at fine resolutions is reduced by nearly an order of magnitude as the reconstruction quadrature is refined. This identifies under-resolved integration as a major contribution in the low-quadrature regime; the remaining plateau is the residual observed at fixed mesh resolution and within the tested quadrature ladder;
3. **positional fidelity does not imply differential fidelity, and the blending contributes additional second-order sensitivity.** The reconstructed position is accurate at the sampled points, yet the curvature error concentrates along the transition regions in which adjacent charts interact. The exact product-rule decomposition exposes terms involving derivatives of the blending weights, and the controlled experiments support their relevance.

The regularity of the parametrisation and its exact single-chart behaviour on the vertex plateaus are formal properties of the DALa construction. The convergence claims, by contrast, remain empirical. Chapter 8 gives an exact decomposition of the blended error and a conjectural comparison with Xu's local polynomial fitting, but it does not prove the stability and approximation bounds required for DALa convergence. What the thesis contributes is therefore a controlled empirical account together with a precise separation between proved properties of the construction, interpretations supported by the experiments, and open approximation questions.

7. Experiments

This chapter reports the experiments that assess the central question set out in Chapter 6: a smooth C^k atlas reconstructs a surface using only *order-zero* (positional) information — it is fitted to the vertex positions of a mesh, with no constraint on normals, curvature, or any higher-order differential quantity — so *does it nonetheless recover those higher-order quantities faithfully?* We answer this on a controlled testbed; the theoretical reading of the curvature results is deferred to the discussion in Chapter 8. A final experiment (Section 7.6) then carries the atlas into a downstream application — functional maps — to test whether its smooth off-vertex evaluation helps there too.

All experiments share a common recipe: a surface for which the exact differential quantities are known in closed form (a torus), several estimators of Gaussian curvature computed on it, and an error measured against the analytic ground truth. The reconstruction itself, and the pointwise evaluation of position and curvature on the resulting atlas, are carried out by an externally supplied C++ library developed by Claudio Mancinelli. The DALa language and atlas construction are described in a manuscript in preparation [19]; Chapter 4 states the mathematical properties used in the analysis, while implementation details are not reproduced here. For the purposes of the experiments, the library is a fixed component which, given a triangle mesh, builds the atlas and returns the reconstructed position $S(p)$ and its Gaussian curvature $K_S(p)$ at any surface point p. Everything around this fixed component — generating the test surfaces, driving the reconstructions, computing the errors, and producing the figures and renders shown below — is orchestrated by a small suite of Python scripts written for this work, so that every result in this chapter is reproducible with a single command.

The code. Each experiment has a single-command driver that runs the relevant C++ binary over the test meshes, reads back its CSV output, and produces the tables and figures: `run_A1.py` (accuracy versus resolution), `run_A2.py` (interior samples), `make_A4_table.py` (analytic versus PN target), `run_quad_ladder.py` (quadrature refinement), `run_gate_comparison.py` (C^k versus soft gate) and `run_dof.py` (chart degree). They share a module of common utilities, `torus_common.py` (paths, the analytic ground truth, and the multi-panel renderer), and a plotting style, `plot_style.py`, that matches the fonts of this document.

Two further scripts render input figures directly, outside the driver loop: `render_meshes.py` (the input tessellations) and `render_curvature_sequence.py` (the curvature-versus-resolution sequence). The functional-maps experiment (Section 7.6) lives in the notebook `Atlas-FunctionalMap-Integration.ipynb`, with helpers in `fmaps_experiment.py`. The C++ code will be available upon publication of the corresponding paper.

The experiments at a glance. Eight experiments are reported, five in this chapter and three in Chapter 8, where they are needed to support the theoretical reading. They share the testbed described in Section 7.1 and differ, by design, in a single ingredient each: Table 1 lists what that ingredient is and what varying it reveals. The split between the two chapters is one of role, not of importance — the experiments of Chapter 7 map *where* the atlas is accurate, those of Chapter 8 isolate *why*.

Table 1. The eight experiments, the single ingredient each of them varies, and what varying it shows. The first five map the behaviour of the atlas; the last three, reported in Chapter 8, test possible explanations.

experiment	what varies	what it shows
A1 (§7.2)	tessellation, resolution	selected atlas fits improve under refinement on irregular tessellations, where the angle-defect error does not decrease
A2 (§7.3)	where the surface is sampled	positional fidelity doesn't imply differential fidelity
A4 (§7.4)	attractor target	the PN surrogate also degrades positional accuracy on irregular meshes
Ladder (§7.5)	quadrature points	the measured error exhibits a sharp quadrature threshold
FMaps (§7.6)	how signals are evaluated	smooth evaluation gives no measurable benefit in this setup
Gate (§8.2.1)	where the gate is applied	the proved C^k gate is less accurate at the tested quadrature
Regularisation (§8.3)	penalty on chart mismatch	damping chart disagreement restores a decreasing error trend
Degree (§8.4)	polynomial degree of the charts	the results point to under-determination of the fit

7.1 Common setup

The test surface. We use the torus of revolution with major radius $R = 1$ and minor radius $r = 0.4$, parametrised by the ambient map

$$\mathbf{x}_{\mathrm{tor}}(u, v) = \Big((R + r\cos u)\cos v,\ (R + r\cos u)\sin v,\ r\sin u\Big),$$

whose Gaussian curvature is known in closed form,

$$K_{\mathrm{true}} = \frac{\cos u}{r\,(R + r\cos u)} = \frac{\rho - R}{r^2\,\rho}, \qquad \rho = \sqrt{x^2 + y^2}.$$

The right-hand expression lets us evaluate the ground truth directly from the Cartesian coordinates of any point on the torus. For a reconstructed point that does not lie exactly on the analytic surface, the same formula is used as a consistent radial extension of the reference curvature; the positional errors reported below quantify the size of this displacement. The subscript "true" identifies this as the analytic reference value of the same Gaussian curvature K used throughout the thesis.

Tessellations and resolutions. To probe robustness we tessellate the same analytic torus in four qualitatively different ways, generated so that every vertex lies exactly on the surface: *regular* (uniform parameter grid), *anisotropic* (triangles elongated along the major circle), *irregular* (vertices jittered off the grid), and *irregular anisotropic* (both). Each tessellation is produced at three resolutions — about 500, 1000 and 5000 vertices — yielding twelve meshes. Isolating the tessellation *type* from the vertex *count* in this way separates sensitivity to mesh quality from the observed error trend under refinement. Figure 14 shows the four input triangulations.

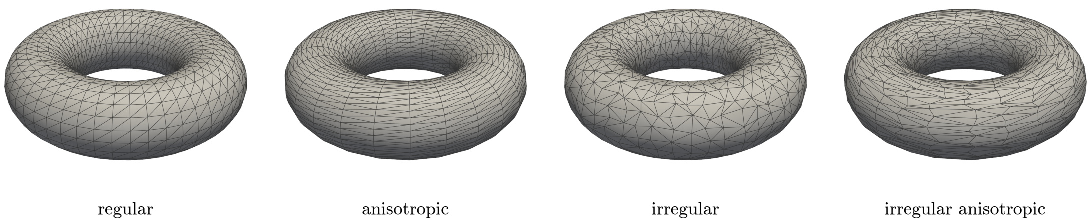

Figure 14. The four input triangulations of the torus (at 1000 vertices): *regular* (uniform grid), *anisotropic* (triangles elongated along the major circle), *irregular* (vertices jittered), and *irregular anisotropic* (both). Every vertex lies exactly on the analytic surface. Produced by `render_meshes.py`.

Curvature estimators. On each mesh we compare six pointwise estimators of the Gaussian curvature at the vertices:

- **discrete** — the classical angle-defect divided by the barycentric vertex area [11]; it uses only the piecewise-linear mesh and serves as the standard baseline;
- **atlas oracle** — the atlas reconstruction fitted to the *analytic* torus target, with curvature read from the first and second fundamental forms of the smooth chart; this is the "best case", in which the reconstruction is driven by the true geometry rather than by a heuristic;
- **strong** — the variational reconstruction with *hard interpolation*, i.e. the charts are constrained to pass through the mesh vertices;
- **weak** — the variational reconstruction with a *pointwise attractor*, i.e. the vertices are matched softly rather than exactly;
- **SA (PN)** and **SA (torus)** — the *surface-attractor* variant, in which the reconstruction is attracted to a target sampled over the *whole* triangle (not only at the vertices), with the target being respectively the PN-triangle heuristic and the analytic torus.

The three variational modes and the two attractor targets are selected through a single option of the reconstruction library; the discrete estimator and the oracle do not depend on the internal quadrature discussed below.

Error measure. For every estimator we denote its Gaussian-curvature output by K_{est} and report the mean absolute error against the ground truth,

$$\mathrm{MAE} = \frac{1}{|V|} \sum_{i \in V} \Big| K_{\text{est}}(v_i) - K_{\text{true}}(v_i) \Big|,$$

averaged over the mesh vertices V (lower is better).

A note on numerical integration. The variational reconstruction assembles its energy and attractor matrices by integrating over each triangle with a fixed quadrature rule. This is the principal numerical integration parameter varied in the experiments, and it becomes the subject of the experiment in Section 7.5. Two rules are used throughout: the standard 46-point rule and a *composite* rule obtained by one midpoint subdivision of the reference triangle ($4 \times 46 = 184$ points); unless stated otherwise, the results below use the composite rule.

7.2 Experiment A1

Goal and protocol. Experiment A1 measures how each estimator behaves as the mesh is refined, and how sensitive it is to the tessellation. The driver `run_A1.py` reconstructs the six estimators on all twelve meshes — with both the 46-point and the composite quadrature — computes the MAE against K_{true}, and produces the plots, tables and surface renders. This delimits where the smooth atlas is genuinely advantageous and where it is not.

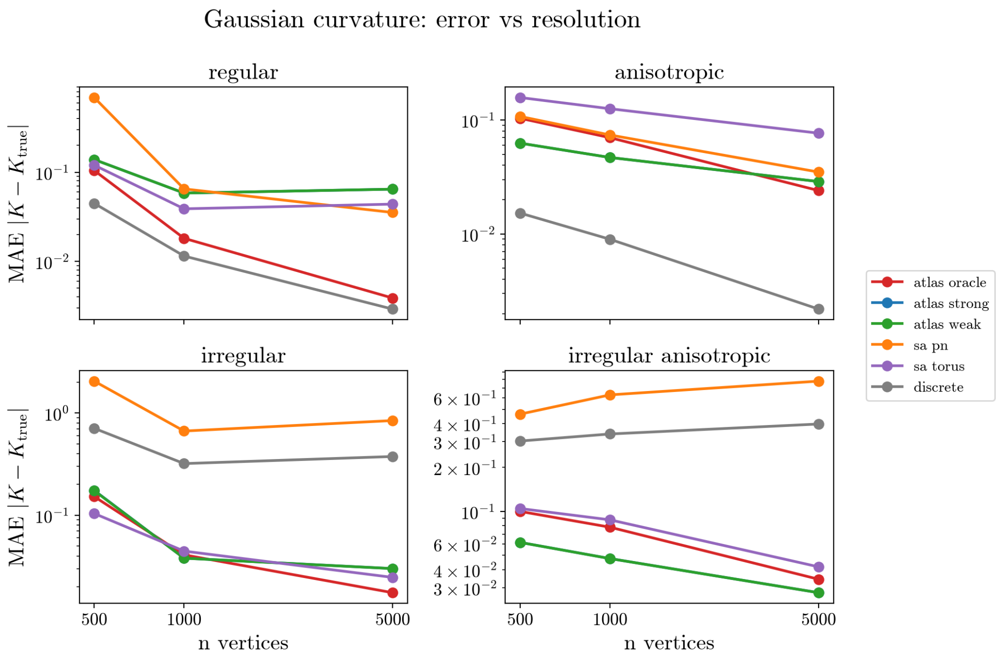

Figure 15. Gaussian-curvature MAE versus vertex count (composite quadrature), one panel per tessellation. On regular/anisotropic meshes the discrete estimator is best; on irregular meshes its error does not decrease, while the strong, weak, and analytically targeted atlas fits retain decreasing errors over the sampled refinements.

Results. Figure 15 plots the MAE against the vertex count, one panel per tessellation. Three facts stand out. First, on *regular* and *anisotropic* meshes the discrete estimator is excellent and its error decreases across the tested refinements (down to ~ 2–3×10^{-3} at 5000 vertices): here the atlas brings no advantage, an honest negative control. Second, and centrally, on *irregular* meshes the discrete estimator *collapses*: its error does not decrease with resolution and can even grow ($\approx 0.71 \to 0.32 \to 0.38$ for the irregular torus), because the poor triangle geometry persists at every scale. The atlas, by contrast, remains robust across these tests and its error decreases under the sampled refinements (the oracle reaching $\approx 1.8 \times 10^{-2}$, the weak reconstruction $\approx 3.0 \times 10^{-2}$ at 5000 vertices) — on the irregular mesh the atlas is an order of magnitude more accurate than the discrete baseline. Third, among the attractor variants the analytically-targeted one (SA torus) retains a decreasing error trend, whereas the PN-triangle one (SA PN) *degrades* on irregular meshes: attracting the surface, over the whole triangle, to a heuristic that itself suffers from bad tessellation amplifies the error rather than reducing it.

The numbers behind Figure 15 are collected in Table 2. Four readings stand out. On the regular and anisotropic tori the discrete estimator is the most accurate and decreases cleanly (0.0447 → 0.0029 and 0.0152 → 0.0022). On the two irregular tori it does not form a decreasing sequence: between the last two resolutions it rises from 0.3203 to 0.3757 and from 0.3387 to 0.3963. The strong and weak atlas estimators improve across the tested refinements and end an order of magnitude below it (0.0301 and 0.0278 at 5000 vertices). Third, the surface attractor with PN-triangle target (`sa_pn`) degenerates on the same irregular meshes (0.8443, 0.7788), whereas the same attractor with the analytic target (`sa_torus`) stays accurate. This comparison points to the PN surrogate, rather than the attractor formulation alone, as the main source of the degradation.

Table 2. Experiment A1 — Gaussian-curvature MAE (composite quadrature), per tessellation and vertex count. Produced by `run_A1.py` (`Exp_A1/mae_table_-composite.csv`). Best value of each row in bold.

tessellation	n	oracle	strong	weak	sa (PN)	sa (torus)	discrete
regular	500	0.1046	0.1387	0.1387	0.6878	0.1200	**0.0447**
	1000	0.0182	0.0584	0.0584	0.0650	0.0389	**0.0115**
	5000	0.0039	0.0646	0.0646	0.0354	0.0439	**0.0029**
anisotropic	500	0.1033	0.0625	0.0625	0.1072	0.1571	**0.0152**
	1000	0.0698	0.0468	0.0468	0.0737	0.1253	**0.0090**
	5000	0.0240	0.0288	0.0288	0.0349	0.0765	**0.0022**
irregular	500	0.1523	0.1750	0.1738	2.0524	**0.1040**	0.7089
	1000	0.0410	**0.0380**	**0.0380**	0.6666	0.0446	0.3203
	5000	**0.0175**	0.0301	0.0301	0.8443	0.0247	0.3757
irregular	500	0.1000	0.0614	**0.0613**	0.4632	0.1046	0.3031
anisotropic	1000	0.0780	0.0476	**0.0476**	0.6273	0.0876	0.3387
	5000	0.0343	0.0278	**0.0278**	0.7788	0.0417	0.3963

The fourth reading is a caveat, and it drives the rest of the chapter. On the *regular* torus the variational estimators do not form a decreasing sequence: the error falls from 0.1387 to 0.0584 and then *rises* to 0.0646 at 5000 vertices. The quadrature and regularisation experiments indicate that this reversal is not intrinsic to the atlas as a representation: refining the quadrature removes it in Section 7.5, and the regularised fit of Section 8.3 yields a decreasing tested sequence. Table 2 should therefore be read as the behaviour of the *default* configuration, not as the best the atlas can do.

The surface renders make the same point visually. Figure 16 colours the torus by the absolute curvature error: on the regular mesh all estimators are uniformly accurate, whereas on the irregular mesh the discrete estimator is a field of bright, per-triangle noise while the atlas panels stay dark and smooth.

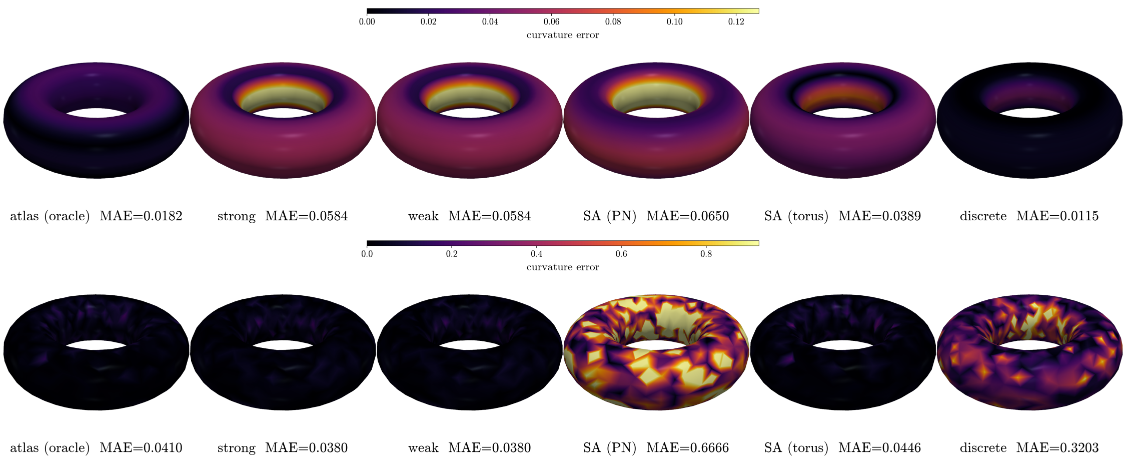

Figure 16. Torus coloured by curvature error $|K_{\text{est}} - K_{\text{true}}|$ (shared colour scale, six estimators), on the regular mesh (top) and the irregular mesh (bottom). The discrete estimator (right panel) degrades sharply on the irregular tessellation; the atlas estimators do not.

Effect of the quadrature. The same run also compares the 46-point and composite rules (Figure 17). With only 46 points the variational estimators appear to stall on the finer meshes; with the composite rule they recover decreasing trends over the tested resolutions. This observation motivates the dedicated study of Section 7.5.

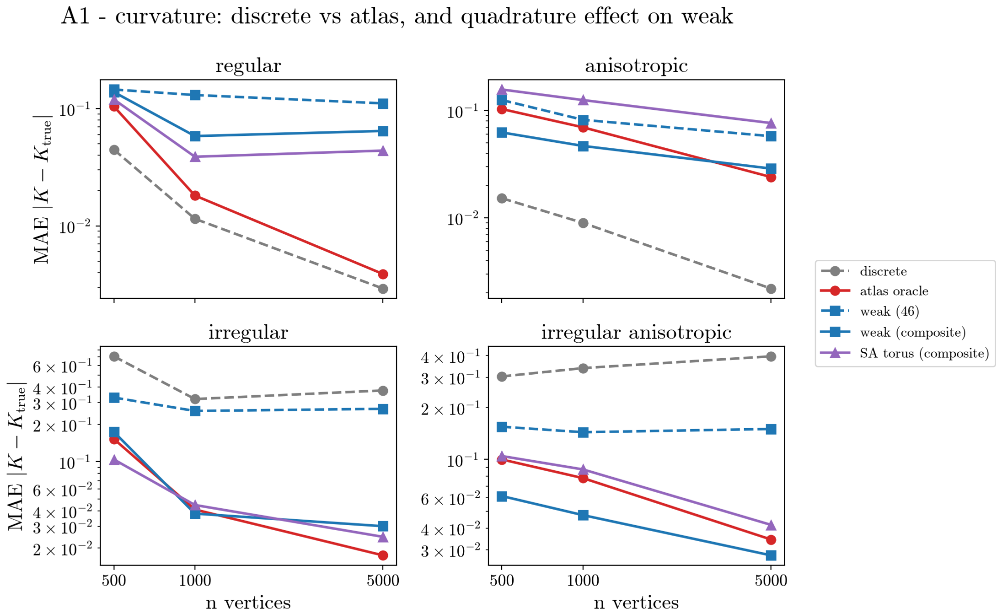

Figure 17. Effect of the reconstruction quadrature: the discrete and oracle estimators are quadrature-independent, whereas the weak reconstruction improves markedly when the 46-point rule is replaced by the composite one — pointing at numerical integration as the limiting factor at fine resolutions.

7.3 Experiment A2

Goal and protocol. A1 measures *how much* the estimators err; A2 asks *where* and *why*. The working hypothesis of the thesis is that positional fidelity does not imply differential fidelity: a reconstruction can match the surface almost perfectly in position while still misjudging its curvature, with the error concentrated at the boundaries between the blending regions of adjacent charts (where several vertex polynomials interact and their jets need not agree). To expose this, A2 samples points in the *interior* of every triangle — not only at the vertices, as in A1 — and, for each sample p and each reconstruction, records

- the *positional error*, the distance of the reconstructed point $S(p)$ from the analytic torus surface;
- the *curvature error* $|K_S(p) - K_{\mathrm{true}}(S(p))|$;
- maxbary, the largest barycentric coordinate of p, a proxy for how deep the sample sits in a blending region: maxbary $= 1$ at a vertex (a single chart dominates), maxbary $= \frac{1}{3}$ at the triangle centroid (the deepest three-chart blend).

The reconstruction uses the composite quadrature so that the effect is not confused with insufficient integration (Section 7.5). The driver `run_A2.py` runs this on the four 1000-vertex tori and aggregates the two errors against maxbary.

Results. The positional error stays tiny across the sampled interiors — of order 10^{-4}, and essentially machine-zero at the vertices, which the strong and oracle reconstructions interpolate exactly (Figure 19). The curvature error, on the other hand, *grows towards the interior blending region*: as maxbary decreases from 1 (vertex) to $\frac{1}{3}$ (centroid) it rises sharply, most dramatically for the oracle. The increase reaches order-one values on several tessellations and includes a much larger mean in the deepest bin of the irregular anisotropic case, where a small number of extreme errors dominate the average (Figure 18). The two fidelities are therefore clearly decoupled in these samples: a near-perfect position can coexist with a large curvature error.

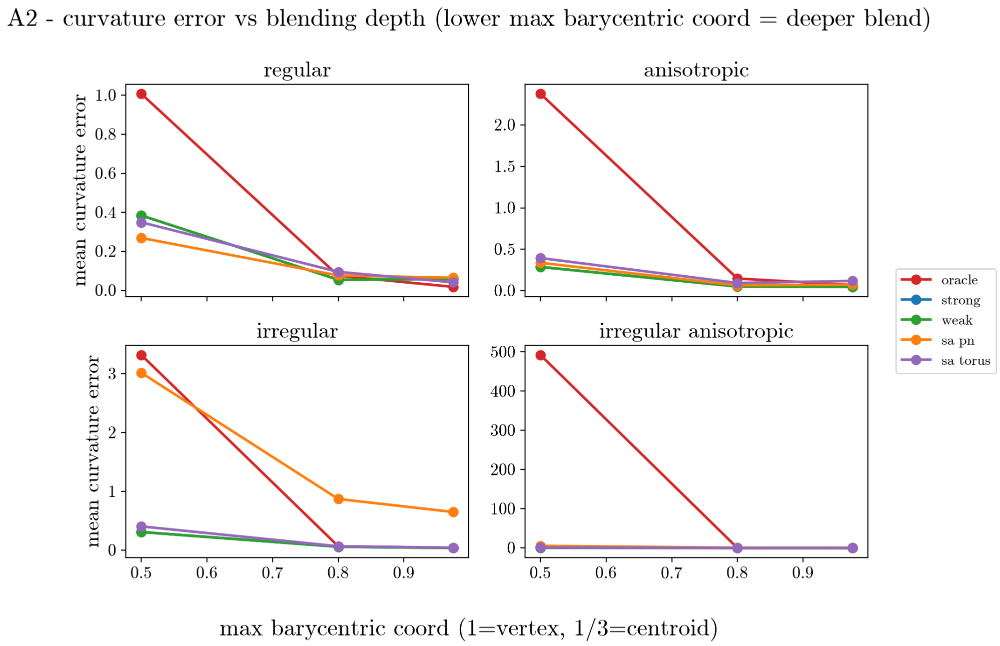

Figure 18. Mean curvature error versus maxbary (axis reversed: deep blend on the left), one panel per tessellation. The error rises towards the interior blending region; the effect is most pronounced for the analytically-targeted oracle and is consistent with additional second-order terms introduced by blending, rather than positional error alone.

Table 3 quantifies the decoupling for the weak reconstruction. The median positional error never exceeds 3.1×10^{-5} (on a torus of tube radius 0.4: a relative error below 0.01%), and even its worst case over all interior samples stays below 10^{-3}. The median curvature error is three orders of magnitude larger. The last column adds a complementary indication: the linear correlation between the two errors is positive but modest (+0.21 to +0.42) — the two errors are related, as one would expect since both grow inside the triangles, but the relation is far too loose for positional accuracy to serve as a proxy. A reconstruction can be right to five decimal places in position and still carry a curvature error of a few percent.

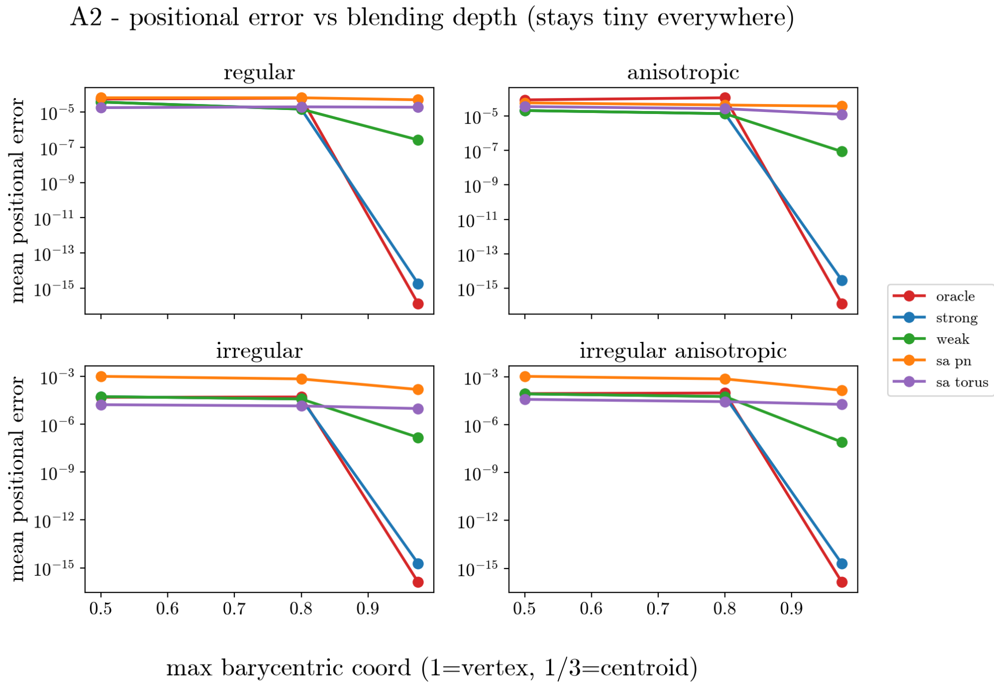

Figure 19. Mean positional error versus maxbary (log scale). Unlike the curvature error, it stays uniformly tiny across the sampled triangle interiors, while the curvature error grows.

Table 3. Experiment A2 — positional and curvature error of the weak reconstruction over the interior samples of the four 1000-vertex tori, and their sample-wise correlation. Produced by `run_A2.py` (`Exp_A2/summary.csv`).

tessellation	positional error		median	corr
	median	max	curv. error	(pos, curv)
regular	1.24×10^{-5}	8.89×10^{-5}	0.0669	+0.42
anisotropic	9.16×10^{-6}	6.91×10^{-5}	0.0423	+0.23
irregular	2.42×10^{-5}	3.74×10^{-4}	0.0623	+0.26
irregular anisotropic	3.07×10^{-5}	9.23×10^{-4}	0.0664	+0.21

The spatial structure of the error is also visible in the interior-point render of Figure 20: the curvature error is spatially organised in bands that follow the chart structure and is concentrated where multiple charts blend.

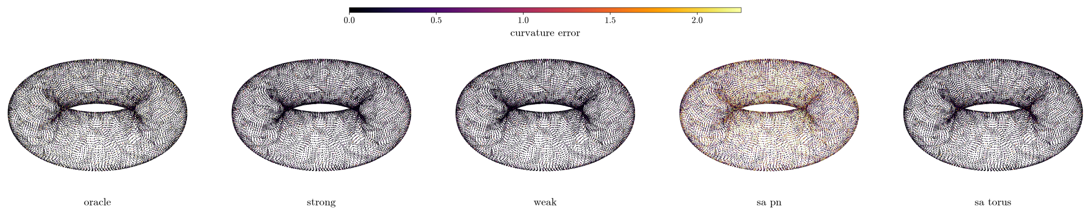

Figure 20. Interior sample points on the irregular torus, coloured by curvature error (five reconstructions). The error is spatially structured and concentrated along the blending boundaries, not uniform.

This is the quantitative evidence for the thesis' leading hypothesis: a faithful positional reconstruction need not yield a faithful reconstruction of the higher-order differential quantities.

The error localises, and refinement hides it. Across the tested refinements, the interior error bands become thinner and fainter, so the curvature field looks progressively smoother even though nonzero errors remain in the shrinking blend regions. Figure 21 shows this directly, evaluating the reconstructed curvature on the interior samples at 500, 1000 and 5000 vertices: the coarse mesh betrays its tessellation in the curvature, the fine one looks smooth. This is the visual counterpart of the mechanism analysed in Chapter 8, and a caution against reading a smooth-looking curvature field as an accurate one.

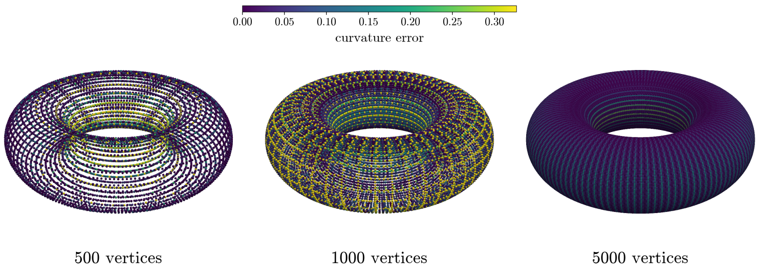

Figure 21. Reconstructed Gaussian-curvature error on the interior samples of the torus at increasing resolution (500 → 1000 → 5000 vertices). The error concentrates in ever thinner bands inside the triangles, so the field looks smoother with refinement while nonzero errors remain in the blending regions. Produced by `render_curvature_sequence.py`.

7.4 Experiment A4

Goal and protocol. Every reconstruction that is not an interpolation needs a *target*: a surface towards which the vertex polynomials are attracted. Two are available. The analytic target (`TorusTarget`) evaluates the true torus at any point of the reference triangle: it is an oracle, unavailable outside a synthetic benchmark. The PN-triangle target (`PNTarget`) is the practical surrogate, a cubic Bézier patch built from the mesh vertices and their normals [15]. A4 asks how much of the atlas' behaviour is due to the fitting machinery and how much to the quality of that target. The comparison is fully controlled and requires no new run: the two surface-attractor estimators of Section 7.2 — `sa_pn` and `sa_torus` — share the fitting mode (`GeometryFitMode::SurfaceAttractor`), the quadrature (composite), the degree, the continuity and the mesh. They differ *only* in the target. The script `make_A4_table.py` extracts the pairs from the A1 and A2 result files.

Results. Table 4 shows a sharp dichotomy. On the two *well-shaped* tessellations (regular, anisotropic) the ratio $\mathrm{MAE}_{\mathrm{PN}}/\mathrm{MAE}_{\mathrm{analytic}}$ stays of order one — between 0.46 and 1.67, with the PN surrogate actually the better of the two at 5000 vertices, where the fit is best resolved. On the two *irregular* tessellations the same ratio explodes: 19.7, 15.0, 34.2 on the irregular torus and up to 18.7 on the irregular anisotropic one. The analytic target keeps improving with resolution (0.1040 → 0.0446 → 0.0247); the PN-targeted sequence is non-monotone (2.0524 → 0.6666 → 0.8443).

The PN target also affects positional accuracy. One might suspect that the PN patch reproduces the surface well and only its second derivatives are wrong.

Table 4. Experiment A4 — curvature MAE of the surface attractor with the PN-triangle target and with the analytic torus target, all else being equal. The last column is their ratio: ≈ 1 on well-shaped meshes, up to $26\times$ on irregular ones. Produced by `make_A4_table.py` (`Exp_A1/a4_pn_vs_analytic.csv`).

tessellation	n	MAE PN	MAE analytic	ratio
regular	500	0.6878	0.1200	5.73
	1000	0.0650	0.0389	1.67
	5000	0.0354	0.0439	0.81
anisotropic	500	0.1072	0.1571	0.68
	1000	0.0737	0.1253	0.59
	5000	0.0349	0.0765	0.46
irregular	500	2.0524	0.1040	19.73
	1000	0.6666	0.0446	14.95
	5000	0.8443	0.0247	34.24
irregular	500	0.4632	0.1046	4.43
anisotropic	1000	0.6273	0.0876	7.16
	5000	0.7788	0.0417	18.66

The interior samples of A2 rule this out. Table 5 reports, on the same meshes, the median positional error of the two attractors: on the well-shaped tessellations both sit in the 10^{-5} range, but on the two irregular ones the PN attractor is a full order of magnitude worse (4.4×10^{-4} against 1.1×10^{-5}). The reconstruction is being pulled towards a surface that is itself off. The comparison therefore points to the surrogate as the main source of the degradation: the PN patch is built from vertex normals, whose estimation is sensitive to badly shaped triangles. The attractor formulation with the analytic target remains among the most accurate estimators in these tests.

The practical consequence matters beyond the torus. On a real mesh the analytic target does not exist, so the atlas must be fitted against a surrogate. A4 shows that the choice of surrogate can dominate the observed degradation as the tessellation worsens. Any robustness claim must therefore state which target and normal-estimation procedure are being used.

Table 5. Median positional and curvature error over the interior samples (1000-vertex tori, data from A2). The PN target degrades the reconstructed *position* too, on precisely the two irregular tessellations. Produced by `make_A4_table.py` (`Exp_A2/a4_interior_pn_vs_analytic.csv`).

	median positional error		median curvature error	
tessellation	**PN**	**analytic**	**PN**	**analytic**
regular	5.36×10^{-5}	1.64×10^{-5}	0.0886	0.0703
anisotropic	4.25×10^{-5}	1.87×10^{-5}	0.0663	0.1140
irregular	4.38×10^{-4}	1.12×10^{-5}	0.8181	0.0716
irregular anisotropic	3.94×10^{-4}	1.97×10^{-5}	0.7808	0.1196

7.5 Quadrature refinement

Goal and protocol. Both A1 (the fine-mesh stall) and the quadrature comparison suggested that the residual error of the variational reconstruction, once the tessellation is good, is strongly affected by the numerical integration of the fitting energy. This experiment tests that interpretation on a single, fixed mesh (the regular torus at 5000 vertices, where the residual was visible) by *refining the quadrature* while keeping everything else constant. The reconstruction library exposes the number of midpoint-refinement levels of the reference-triangle rule as a run-time parameter, giving the ladder $46 \to 184 \to 736 \to 2944$ points (a $4\times$ increase per level). The driver `run_quad_ladder.py` reconstructs the strong and weak estimators at each level; the discrete and oracle estimators, being quadrature-independent, provide constant reference lines.

Results. The curvature error falls from 0.1110 to 0.0646, then drops sharply to 0.0070 at 736 points — a single refinement worth a factor of 9 — and finally saturates at 0.0052 (Figure 22). Two conclusions follow. First, at low quadrature the dominant part of the residual is an integration error: the stall observed on the regular mesh in Section 7.2 is removed by refining the rule alone, without touching anything else. Second, the gain is concentrated in *one* step. Neither the first refinement ($1.7\times$) nor the third ($1.35\times$) matters much; the second, from 184 to 736 points, is where the error drops by an order of magnitude. This pattern is consistent with a sampling threshold rather than a smooth asymptotic rate. Section 8.2 relates it to the thin bands in which the blending weights vary sharply, while Section 8.2.1 tests how the gate changes their numerical effect.
Unlike the case of a smoother gate (Section 8.2.1), the plateau at 0.0052 does *not* overtake the discrete baseline (0.0029), and stays just above the oracle (0.0039). On

the regular torus, with the default fit and this mesh, refining the quadrature alone is therefore not enough to beat the angle defect: it recovers a decreasing tested sequence but not superiority. The regularised fit of Section 8.3 reaches 0.0065 using only the 184-point rule — essentially the accuracy of the 2944-point ladder at a sixteenth of the quadrature — which makes it, and not further refinement, the practical lever.

Table 6 adds the cost side of the trade-off. The reconstruction time grows more slowly than the number of quadrature points — roughly 2–4× per level, since part of the cost is independent of the rule — so the largest observed gain, the refinement worth a factor of 9 in accuracy, costs only about 2.8× the time. The third refinement is the bad deal: 3.7× the time for a 35% error reduction. A1 and A2 use the 184-point composite rule as a tractable default. Within this four-level experiment, the 736-point rule gives the best observed accuracy–cost compromise for this mesh.

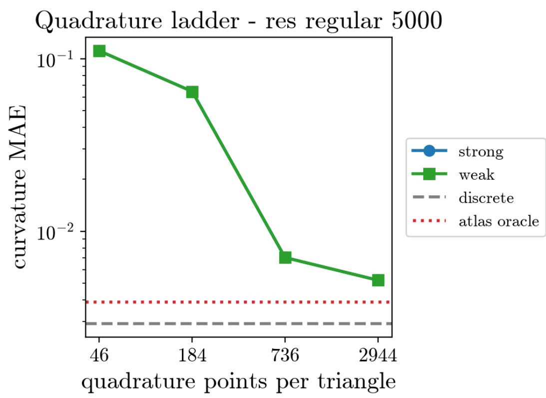

Figure 22. Curvature MAE versus quadrature points at fixed mesh (regular torus, 5000 vertices). The gain is concentrated in a single step, from 184 to 736 points, where the error drops by a factor of 9; the pattern is consistent with under-resolution below this level, while the error then approaches 0.0052, just above the oracle and still above the discrete reference.

7.6 Functional Maps experiment

The curvature experiments probe a *geometric* second-order quantity — derivatives of the charts themselves. This last experiment asks a complementary question about the atlas' other distinctive ability: evaluating a *scalar field* continuously at any point of a triangle, rather than only at the vertices. Does that continuous, off-vertex evaluation help a pipeline that consumes scalar fields? We test it on functional maps [3], the downstream application of Part II, which are assembled entirely from spectral signals (Laplace–Beltrami eigenfunctions and heat-kernel

Table 6. Quadrature ladder on the regular torus at 5000 vertices: curvature MAE and reconstruction time per refinement level. The largest observed gain is the second refinement (184 $\rightarrow$ 736 points), worth a factor of 9 in accuracy for 2.8$\times$ the time. The discrete (0.0029) and oracle (0.0039) references do not depend on the quadrature and are never overtaken. Produced by `run_quad_ladder.py` (`Exp_quad_ladder/ladder_table.csv`).

level	points/triangle	MAE strong	MAE weak	t_{strong} [s]	t_{weak} [s]
0	46	0.1110	0.1110	19.6	9.8
1	184	0.0646	0.0646	30.5	19.0
2	736	0.0070	0.0070	61.3	52.7
3	2944	**0.0052**	**0.0052**	203.4	193.5

descriptors) sampled and integrated over the surface.

Protocol: one pipeline, three ways to evaluate the signals. We compute the same functional map between a pair of near-isometric FAUST shapes [14] (the registrations `tr_reg_000` and `tr_reg_004`) in three ways, changing *only* how the per-vertex signals are turned into values at $N = 5000$ surface sample points:

- **discrete** — the signals are used at the mesh vertices and projected with the mass-matrix inner product; this is the standard vertex-based baseline;
- **linear** — the signals are evaluated at the samples by barycentric (piecewise-linear) interpolation within each face and projected by the area-weighted quadrature of Section 7.1 of Part I;
- **atlas** — the *same* per-vertex signals are evaluated at the *same* samples through the smooth atlas instead of by linear interpolation.

Everything downstream is shared across the three: the truncation to $k_{\text{fmap}} = 30$ eigenfunctions, the least-squares solve for the functional map matrix $\mathbf{C}$, its conversion to a point-to-point map, and the evaluation of that map (geodesic and Euclidean error against the ground-truth correspondence, coverage, and the “diagonal score” that measures how close $\mathbf{C}$ is to the diagonal expected of a near-isometry). The comparison is therefore controlled: any difference between *linear* and *atlas* is due to the reconstruction of the signals alone.

How the atlas evaluation is wired in. The atlas lives in the C++ library and the functional-maps pipeline in a Python notebook, so the two are bridged by an *export–compute–reload* exchange, implemented in `fmaps_experiment.py` and driven from `Atlas-FunctionalMap-Integration.ipynb`. The notebook exports

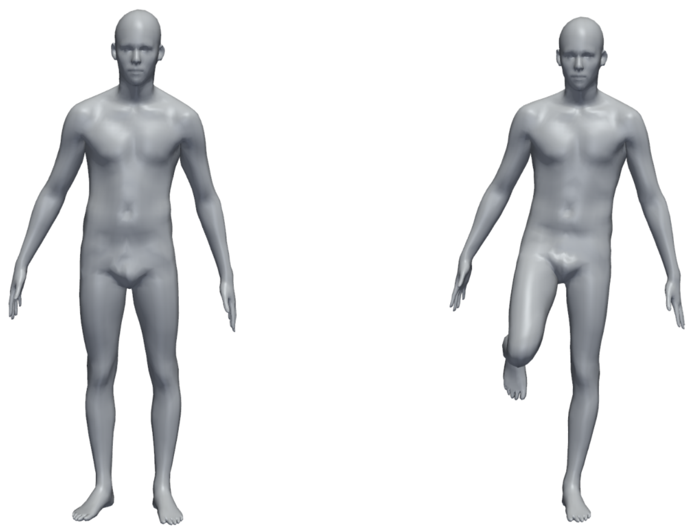

source (`tr_reg_000`) target (`tr_reg_004`)

Figure 23. The two near-isometric FAUST registrations used in the experiment: `tr_reg_000` (source) and `tr_reg_004` (target), the same human in two poses. The functional map between them is expected to be close to diagonal.

the mesh (in the vertex ordering it uses), the barycentric coordinates of the N sample points, and the per-vertex signals; the C++ program `fmaps_test`, run in its `reload` mode, builds the atlas over that mesh and, treating each signal column as vertex data, reconstructs it as a smooth scalar field and *evaluates it at the exported sample points*; the notebook reloads the resulting columns and feeds them into the shared pipeline in place of the linearly-interpolated ones. The reconstruction is the same C^k atlas used throughout Part I ($k_{\text{cont}} = 2$, vertex polynomials of degree three), so the experiment exercises the very object the rest of the thesis analyses.

Result. Table 7 reports the three paths. The atlas evaluation is numerically almost indistinguishable from the linear one in this setup — the two agree to three decimals on every metric, differing only in the fourth (diagonal score 0.86180 versus 0.86189; geodesic error 0.016617 versus 0.016606) — while the plain discrete baseline is in fact marginally the best. Replacing linear interpolation of the signals by their smooth atlas reconstruction does not materially change the recovered map, while it carries a steep price: the atlas evaluation of the signals takes about eight minutes, against a fraction of a second for the linear one (Table 7, last column).

The same conclusion is visible without reading a single number. Figure 24 shows the three recovered maps **C** as heat maps: all three are dominated by a near-diagonal

Table 7. Functional map between two near-isometric FAUST shapes, computed three ways ($N = 5000$ samples, $k_{\text{fmap}} = 30$), changing only how the signals are evaluated. Higher diagonal score and coverage are better; lower geodesic and Euclidean error are better. The atlas path matches the linear one to the fourth decimal, but its signal evaluation — the C++ atlas reconstruction of every signal column at the N samples — costs about 469 s ($\approx$ 8 minutes) against well under a second for either baseline. Produced by `Atlas-FunctionalMap-Integration.ipynb` (`ExpFMAPS/scores.json`, `params.json`).

path	diag. score	geodesic err.	Euclidean err.	coverage	time [s]
discrete	**0.8680**	**0.0142**	**0.0147**	**0.6718**	< 1
linear	0.8619	0.0166	0.0169	0.6310	< 1
atlas	0.8618	0.0166	0.0169	0.6314	469

band — as expected for a near-isometry between the two FAUST shapes — and the linear and atlas panels are visually identical, down to the faint off-diagonal structure. The smooth atlas reconstruction of the signals reproduces the linearly-interpolated map entry by entry.

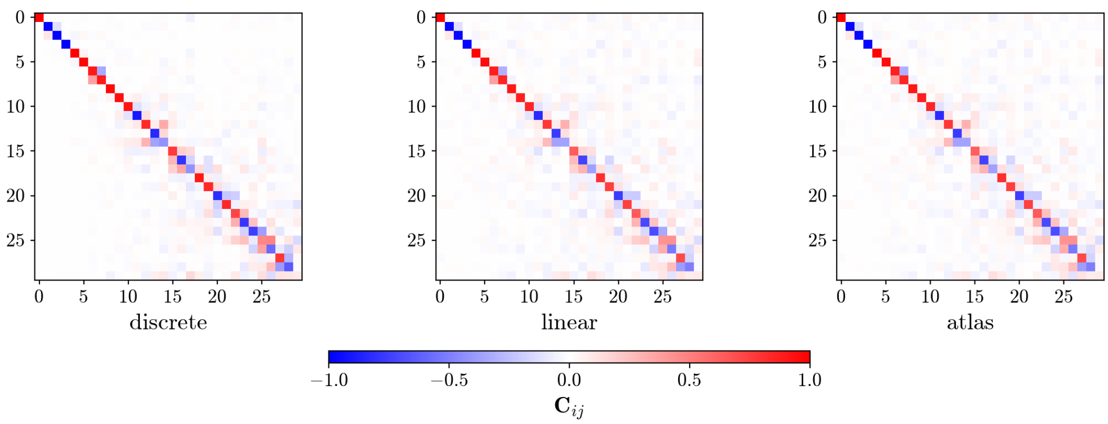

Figure 24. The 30×30 functional maps $\mathbf{C}$ recovered along the three paths (same setup as Table 7). All three are near-diagonal; the linear and atlas panels are visually indistinguishable, consistently with the near-equality of the numerical metrics. Produced from `ExpFMAPS/fmaps.npz`.

Why the smooth evaluation gives no measurable benefit here — in contrast to the selected curvature experiments, where it improves the estimate — is taken up in Section 8.5.

8. Discussion

Chapter 7 is empirical: it reports what the atlas does to the Gaussian curvature of a torus. Here we conjecture how the observed behaviour may arise and provide a sketch of a possible proof strategy. The comparison with classical local polynomial fitting explains why decreasing vertex errors are plausible, but it does not establish convergence of the DALa reconstruction. We then use an exact product-rule decomposition to interpret the interior artefact in terms of the blending weights, and finally discuss the downstream functional-maps experiment.

8.1 A conjectural convergence argument

The positive evidence to be explained is deliberately limited: selected DALa configurations show decreasing vertex-curvature errors along the tested refinements, including irregular tessellations for which the angle defect does not improve. This finite experimental trend is not, by itself, an asymptotic result. A useful route towards understanding it is to compare the local problem seen at a vertex with the polynomial-fitting framework of Xu [16].

The local viewpoint. At a mesh vertex, Equation (4.3) already states that the DALa blend reduces to the corresponding vertex patch on an open plateau. Thus the pointwise vertex curvature is determined by the derivatives of that local patch, without additional derivative terms from the partition of unity. This is a direct consequence of the construction described in Chapter 4, not a new approximation result. The non-trivial question is instead whether the coefficients computed for the patch approximate the jet of the underlying smooth surface as the mesh is refined.

Comparison with Xu. Xu studies a local least-squares polynomial fitted to samples of the piecewise-linear input surface, expressed in a planar coordinate system around a vertex [16]. For a quadratic fit and a suitably well-posed rescaled node set, the analysis gives convergent estimates of first- and second-order differential quantities. The fit is used pointwise at the vertex: no globally blended surface is constructed, and its coefficients are determined directly by the sampled local geometry.

The DALa vertex construction shares two relevant ingredients with this setting. First, it also works on a locally flattened vertex neighbourhood, with the conformal construction described in Section 4.1. Second, although DALa represents its patches in the triangular Bernstein basis, the degree-d triangular Bernstein polynomials form a basis of the complete polynomial space $\mathcal{P}_d$. The choice between this basis and the standard monomial basis used in a conventional polynomial fit is therefore not the essential distinction at fixed degree.

The essential distinction is how the geometry and the coefficients enter the fit. Xu estimates the coefficients by local least squares against samples of the piecewise-linear geometry. DALa instead obtains them from a globally assembled variational problem, subject to interpolation or attractor constraints. In the surface-attractor mode the target may be a PN triangle: a cubic Bézier proxy built from mesh vertices and normals, chosen as a practical smooth surrogate rather than as the piecewise-linear input itself. In principle that target could be replaced by samples of the piecewise-linear surface, which would make the local geometry proxy closer to Xu's; the variational coupling and numerical quadrature would nevertheless remain different.

Conjecture and proof sketch. The experiments suggest the following cautious conjecture: when the local coefficient estimation is stable and consistent with the sampled geometry, the pointwise DALa vertex derivatives may inherit a convergence mechanism analogous to Xu's local fitting result. A proof could compare the computed DALa patch on a rescaled vertex domain with an appropriate local least-squares polynomial, then bound the difference between their coefficients and transfer that bound to the required derivatives. Such a comparison would have to control the chosen geometry proxy, the variational energies, global coupling, and quadrature uniformly under refinement. None of those bounds is established here, so no formal rate is claimed for the current DALa solves.

This sketch also marks the limitation of the comparison. Xu's stated result is quadratic and stops at second-order quantities. The degree-three experiments do not justify extrapolating its rate, and increasing the degree does not automatically improve a fit whose coefficients are insufficiently constrained. A broader result, independent of a particular polynomial degree and applicable to differential quantities of arbitrary order, remains a question for future work.

8.2 Blending weights and localised artefacts

The weights enter the curvature through two distinct channels, and it is worth separating them because they act at different stages; the first concerns the unresolved fitting problem of Section 8.1.

The first channel is the *fit*. As noted above, the gated basis is used to assemble the variational energy, so the choice of gate changes the fitted vertex charts and hence the vertex curvature, even though the weights play no role in evaluating it on the plateau. This is precisely why a complete convergence argument will have to control the dependence of the fitted vertex charts on the gate and on the global variational coupling. The next section measures that dependence directly.

8.2.1 Provable regularity versus quadrature demand

There is a natural experiment for the claim above, and it is not artificial: the construction admits two ways of applying the gate. Applying it to the blending *weight* is the construction for which C^k compatibility is proved; applying it to the barycentric coordinate — the variant used in an earlier version of the library — produces a visually similar partition of unity with gentler transitions, but carries no such guarantee. The two differ in nothing else: same ε, same charts, same composite quadrature, same estimators. Table 8 compares them.

The controls behave consistently with this distinction: the oracle and the discrete estimator, which do not pass through the blended fit, are identical to four decimals on every tessellation (ratio 1.00). Every variational estimator, by contrast, moves — and it moves in the direction that the error decomposition anticipates. The C^k-compatible gate is the *less* accurate of the two, by up to 2.7× on the regular torus. This is consistent with its sharper transition producing larger weight derivatives at the employed quadrature, and hence amplifying terms (b) and (c); the experiment measures this dependence but does not by itself prove the causal scaling.

Two details make the reading sharper. First, the penalty is not uniform: it is largest on the regular torus (0.37), where Section 8.4 suggests that the fit is especially under-determined, and it becomes negligible on the irregular one (1.02, 1.01). This is consistent with chart disagreements dominating the product $\nabla w_i \cdot \nabla e_i$ there. Second, the two gates cost the same time to reconstruct (3.5 s in both cases): the difference is entirely one of accuracy at fixed quadrature, not of computational effort.

This is the observed trade-off in its cleanest form. The gentler gate mitigates the abrupt weight variation and buys accuracy at low quadrature, but the surface it produces is not covered by the C^k compatibility proof. At the tested quadrature, the C^k gate keeps that guarantee but produces the larger errors reported above. Together with the sharp drop in Section 7.5, this supports the interpretation that narrow transition bands are under-resolved at low quadrature; it does not establish a universal cost of C^k compatibility.

Table 8. Vertex-curvature MAE with the C^k-compatible gate (applied to the blending weight) and with the softer legacy gate (applied to the barycentric coordinate), all else equal: same ε, same composite quadrature, 1000-vertex tori. The ratio is $\mathrm{MAE}_{\mathrm{soft}}/\mathrm{MAE}_{C^k}$, so a value below one means the soft gate is the more accurate. The oracle and discrete estimators do not pass through the blended fit and are unchanged (ratio 1.00), a clean control. The measured penalty is largest on the regular torus and negligible on the irregular one. Reconstruction time is the same for both ($\approx 3.5\,\mathrm{s}$). Produced by `run_gate_comparison.py` (`Exp_gate/gate_table.csv`).

		MAE		
tessellation	**estimator**	**C^k gate**	**soft gate**	**ratio**
regular	strong / weak	0.0584	**0.0216**	0.37
	sa (PN)	0.0650	**0.0330**	0.51
	sa (torus)	0.0389	**0.0219**	0.56
anisotropic	strong / weak	0.0468	**0.0401**	0.86
	sa (PN)	0.0737	**0.0382**	0.52
irregular	strong / weak	**0.0380**	0.0389	1.02
	sa (torus)	**0.0446**	0.0452	1.01
irregular an.	strong / weak	0.0476	**0.0418**	0.88
	sa (torus)	0.0876	**0.0719**	0.82
(control) regular	oracle	0.0182	0.0182	1.00
	discrete	0.0115	0.0115	1.00

Figure 25 makes the same point visually, and adds one observation the table cannot convey. Reconstructed with the two gates and coloured by curvature error, the torus shows the *same* spatial pattern in both cases — the error concentrates on the inner rim, where the surface bends most — but the pattern is far brighter under the C^k gate. What changes is the amplitude of the artefact, not its location: the two reconstructions agree on *where* the curvature is hard to recover, and disagree only on how badly. This is what one expects if the error is the product $\nabla w_i \cdot \nabla e_i$

of the decomposition, with the gate rescaling the first factor while the second, which depends on the geometry, stays put. The rendered positions remain visually indistinguishable between the two.

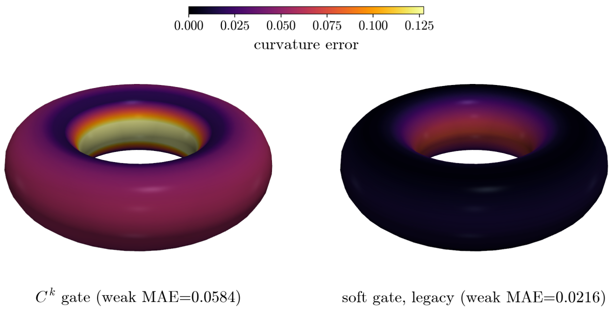

Figure 25. Torus (1000 vertices) coloured by the curvature error of the weak reconstruction, built with the C^k-compatible gate (left) and with the softer legacy gate (right), all else equal. The positional reconstruction is visually the same and the error has the same spatial signature in both; only its amplitude changes, roughly by the factor of 2.7 of Table 8. This is consistent with the weight–disagreement coupling in the product-rule decomposition. Produced by `run_gate_comparison.py`.

The second channel is the *evaluation*, and it is what Section 7.3 exposes. Away from the vertex plateaus, several charts are active at once and the curvature is read from the blended parametrisation itself, so the weights enter directly. All derivatives must be taken in a common triangle coordinate system. We therefore write

$$\widehat{S}_i(u) = S_i(\varphi_{T,i}(u)), \qquad S_T(u) = \sum_i w_i(u)\,\widehat{S}_i(u), \qquad \sum_i w_i(u) \equiv 1.$$

Let $S_T^\star$ be the true surface expressed in the same triangle coordinates; the subscript has the same local-coordinate meaning as in S_T. Define $e_i = \widehat{S}_i - S_T^\star$. Because the weights sum to one,

$$S_T \;=\; S_T^\star + \sum_i w_i\, e_i,$$

so the entire reconstruction error is $\sum_i w_i e_i$. Differentiating twice and using $\sum_i w_i \equiv 1$ (hence $\sum_i \nabla w_i \equiv 0$ and $\sum_i \nabla^2 w_i \equiv 0$),

$$\nabla^2 S_T - \nabla^2 S_T^\star \;=\; \underbrace{\sum_i w_i\,\nabla^2 e_i}_{\text{(a)}} \;+\; \underbrace{2\sum_i \nabla w_i \otimes \nabla e_i}_{\text{(b)}} \;+\; \underbrace{\sum_i e_i\,\nabla^2 w_i}_{\text{(c)}}\,.$$

Term (a) is the direct second-order error of the pulled-back patches. Terms (b) and (c) couple their disagreements with derivatives of the weights and isolate the additional sensitivity introduced by blending.

What the decomposition suggests. The identity above does not prove convergence, but it makes the missing estimates explicit. If a family of local patch errors and their first two derivatives were known to decrease at rates compatible with the growth of the first and second derivatives of the weights, then all three terms could decrease together. A curvature estimate would additionally require the parametrisations to remain uniformly regular, so that the denominator in the curvature formula does not degenerate. Establishing these hypotheses for the actual variational and quadrature solutions is part of the open problem, rather than a consequence of the algebraic decomposition.

Two qualitative consequences are already exact. First, throughout a vertex plateau one weight is identically one and the others, together with all their derivatives, are identically zero. Terms (b) and (c) therefore vanish on an open neighbourhood, not merely in the limit at the vertex. This is consistent with the near-zero error observed as maxbary $\to 1$ in Section 7.3. Second, away from the plateaus the artefact depends on the product of the weight derivatives and the chart disagreements. The experiments indicate that the latter generally decrease under refinement, but this trend is an observation rather than an assumption-free theorem.

Parametric versus physical scaling. The parameter ε in Equation (4.2) is a threshold in the triangle domain. It should not be identified with a physical bandwidth εh, and the present analysis does not justify a particular dependence of the derivative bounds on ε. For fixed ε, the passage from parametric to physical derivatives is governed by the triangle map F_T with Jacobian J_T:

$$\nabla_x w = J_T^{-\top} \nabla_u w, \qquad D_x^2 w = J_T^{-\top} (D_u^2 w) J_T^{-1}$$

for an affine triangle map. Thus the natural bounds involve $\|J_T^{-1}\|$ and, on an anisotropic triangle, its smallest singular value or minimum altitude rather than a single isotropic diameter h. Any future proof based on isotropic $O(h^{-r})$ scaling would therefore need a shape-controlled family; for anisotropic sequences it must retain the corresponding Jacobian bound.

This scaling sharpens the interpretation of the gate experiment. The C^k-compatible construction has a sharper parametric transition than the legacy variant, and at fixed quadrature it produces the larger errors reported above. This is consistent with larger $C_{w,r}(\varepsilon)$, but neither the experiment nor the construction's regularity guarantee determines the dependence of those constants on ε.

The same distinction corrects the anisotropy conjecture. A gate transition has a fixed description in the reference triangle, while its physical width and derivatives are transformed by J_T. On an elongated triangle the small singular value of J_T both compresses physical bands and amplifies transverse derivatives; a fixed reference quadrature is transformed by the same anisotropic map. A high-derivative region can therefore be under-resolved precisely where e_i and ∇e_i must stay small for terms (b) and (c) to remain controlled. If this reading is correct, quadrature adapted to the Jacobian or to the weight derivatives should help the anisotropic meshes disproportionately; verifying this is left to future work.

8.3 Regularising the fit

The decomposition of Section 8.2 presents the additional blending terms as couplings between derivatives of the weights and the chart disagreements e_i and ∇e_i. The tested softer gate reduces the first factor but is not covered by the C^k proof, while resolving the narrow bands with more quadrature is expensive. The next experiment therefore acts on the second factor.

The reconstruction library offers a direct handle on it. Besides the usual energy, the fit can be assembled with two additional terms that penalise the mismatch, at each vertex, between the values of the polynomials meeting there and between their *first jets*. The intent is exactly the one the formula suggests: charts that agree in value and in first derivative at the vertices have small e_i and small ∇e_i, so terms (b) and (c) are damped at the source rather than resolved after the fact. We ran the full set of meshes twice, with and without these terms, at the same 184-point composite quadrature; Table 9 reports the outcome.

The controls are again clean: the oracle and the discrete estimator are unchanged. The variational estimators improve almost everywhere, and the improvement grows with resolution — 10× for the strong fit on the regular torus at 5000 vertices, 3.1× on the irregular one, 2.3× on the irregular anisotropic. The surface attractor with the PN target, the weakest estimator of the whole study, gains between 2.2× and 8.3×, though it remains poor in absolute terms on the irregular meshes: the regulariser cannot repair a target that is itself wrong. At 500 vertices the picture is mixed, and on the anisotropic torus the regularised fit is slightly *worse* — consistent with the counting argument of Section 8.4, since adding penalty terms to an already under-constrained system does not necessarily help.

Two consequences deserve emphasis. First, the regularised fit *restores a monotone decrease across the tested regular-torus refinements*: the sequence $0.0865 \to 0.0213 \to 0.0065$ is monotone, where the unregularised one stalls and reverses $(0.1387 \to 0.0584 \to 0.0646)$. Thus the anomaly flagged in Section 7.2 can be removed by changing the fit while keeping the atlas representation fixed; in this experiment it belongs to the unregularised configuration at the chosen quadrature.

Second, and more practically, the regularised fit at 184 quadrature points reaches 0.0065, essentially matching the 2944-point ladder (0.0052) at a sixteenth of the integration cost. Damping the chart disagreements is far cheaper in this comparison than resolving the weight-transition bands by brute-force sampling, consistently with the product-rule interpretation.

Table 9. Curvature MAE without and with the value-and-first-jet regularisation of the fit, all else equal (184-point composite quadrature). Ratio is $\mathrm{MAE}_{\mathrm{plain}}/\mathrm{MAE}_{\mathrm{reg}}$: above one means the regularisation helps. The improvement grows with resolution and restores a decreasing trend on the tested regular-torus refinements. The oracle and discrete estimators are unchanged. Produced by the `torus_curvature_reg` binary (`gcurv_res_*_reg.csv`).

			strong			**sa (PN)**	
tessellation	n	**plain**	**reg.**	**ratio**	**plain**	**reg.**	**ratio**
regular	500	0.1387	0.0865	1.60	0.6878	0.0826	8.33
	1000	0.0584	0.0213	2.75	0.0650	0.0214	3.03
	5000	0.0646	**0.0065**	10.01	0.0354	**0.0048**	7.35
anisotropic	500	0.0625	0.0787	0.79	0.1072	0.0745	1.44
	1000	0.0468	0.0529	0.88	0.0737	0.0480	1.54
	5000	0.0288	0.0182	1.58	0.0349	0.0145	2.40
irregular	500	0.1750	0.2047	0.85	2.0524	0.7294	2.81
	1000	0.0380	0.0290	1.31	0.6666	0.2499	2.67
	5000	0.0301	0.0096	3.12	0.8443	0.3537	2.39
irregular anisotropic	500	0.0614	0.0602	1.02	0.4632	0.2066	2.24
	1000	0.0476	0.0408	1.17	0.6273	0.2749	2.28
	5000	0.0278	0.0119	2.33	0.7788	0.3561	2.19

8.4 Degrees of freedom versus effective constraints

One observation of Section 7.2 has so far been left without an explanation: on the *regular* tessellation — the easiest of the four — the variational estimators plateau at fine resolution and are beaten by the plain angle defect, while on the irregular tessellations they win by an order of magnitude. Section 7.5 showed that refining the quadrature removes the plateau, which identifies the symptom but not the cause: *why* should the reconstruction on the simplest mesh be the one that needs more integration points?

The explanation we propose is a counting argument. The variational fit determines the chart coefficients from positional information sampled at the quadrature nodes, so the *effective* constraints on the coefficients are not the mesh vertices but the quadrature points at which the fitting energy is assembled. Whenever the degrees of freedom of the charts approach or exceed what those samples can pin down, the leftover degrees of freedom are decided by the energy rather than by the data — invisibly at order zero, since many coefficient configurations reproduce the positions equally well, but visibly at second order, where the unconstrained directions are free to oscillate. The regular torus is the setting in which the experiments show the clearest symptoms of such an imbalance. Its simple geometry and highly symmetric parametric domains suggest that the positional samples may carry less independent information per quadrature point than on an irregular mesh; this interpretation is tested below, not assumed as a theorem.

This hypothesis makes a sharp prediction: *raising* the polynomial degree of the charts (more degrees of freedom, same constraints) should degrade the variational estimators, while an estimator that does not pass through the variational fit should, if anything, improve. We tested it with a dedicated experiment: a separate binary, `torus_curvature_d5k3` (the reconstruction compiled with degree-5 charts and $k_{\text{cont}} = 3$), run over the same meshes and driven by its own script `run_dof.py`, with everything else held fixed (composite quadrature). Table 10 reports the outcome on the regular mesh — the worked case of the counting argument above — and it separates the estimators in the direction predicted by the hypothesis.

Two features of Table 10 carry the argument. First, the oracle — which evaluates the analytic target directly and involves no variational solve — *benefits* from the higher degree at every resolution: when the information to constrain them is available, the extra degrees of freedom pay off. Second, and more tellingly, on this mesh even the surface attractor with the *analytic* torus target degrades by a factor of 3, although its target is exact and is sampled over the whole triangle rather than at the vertices alone. On the regular torus this result points to the *variational solve*, rather than target quality alone, as the bottleneck; its effective constraints are tied to the quadrature samples. We also tested the complementary direction: lowering the blending continuity to $k_{\text{cont}} = 2$, which changes the weights but not the number of chart coefficients, leaves the picture essentially unchanged (the errors move by less than 15%) — consistent with the fact that k_{cont} does not enter the counting.

The prediction is not special to the regular mesh, but the tessellation does modulate it. Table 11 repeats the degree-3 versus degree-5 comparison on all four tessellations and reports, for each estimator, the ratio of the two errors at the finest resolution (5000 vertices). Two groups separate cleanly. The oracle sits at or below one everywhere (0.85–1.01): with a target that constrains them, the extra degrees of

Table 10. Gaussian-curvature MAE on the regular torus with degree-3 (the setting of Chapter 7) and degree-5 charts, all else equal (composite quadrature, $k_{\text{cont}} = 3$). The oracle, which does not pass through the variational fit, improves at every resolution; every variational estimator degrades sharply, including the surface attractor with the *analytic* target. The discrete estimator is independent of the reconstruction and acts as a control. Produced by the `torus_curvature_-d5k3` binary (`gcurv_res_regular_*_d5k3.csv`). Reconstruction time grows by a factor of 4–6 with the higher degree (e.g. strong at 5000: 26 s → 174 s).

		MAE		**time [s]**	
n	**estimator**	**degree 3**	**degree 5**	**degree 3**	**degree 5**
500	oracle	0.1046	**0.0436**	0.0	0.1
1000	oracle	0.0182	**0.0147**	0.1	0.1
5000	oracle	0.0039	**0.0036**	0.3	0.6
1000	weak	0.0584	0.1931	2.8	9.6
5000	weak	0.0646	0.2397	17.3	77.9
5000	strong	0.0646	0.2397	26.4	173.7
5000	sa (PN)	0.0354	0.1546	32.1	111.0
5000	sa (torus)	0.0439	0.1309	31.8	110.6
all	discrete	unchanged		—	

freedom pay off. The purely positional fits sit well above one everywhere (2.6–4.2): with only vertex positions to go on, they do not.

Between the two lies the surface attractor with the analytic target, and its behaviour sharpens the argument. On the anisotropic and irregular-anisotropic tori it sits at 0.86 and 0.98, behaving like the oracle. On the *regular* torus it degrades by 2.98, joining the positional fits. Since the estimator and target are held fixed while the tessellation changes, the observation is consistent with the counting hypothesis: the symmetric regular grid may provide fewer independent effective constraints through the chosen quadrature. Establishing the rank or conditioning of those constraints would be needed to turn this interpretation into a proof.

This reading offers a common explanation for the two anomalies of the regular mesh. Refining the quadrature (Section 7.5) adds effective constraints at fixed degrees of freedom; raising the degree adds degrees of freedom at fixed constraints. The two knobs move the same ratio in opposite directions, and the experiments respond accordingly. The results are therefore consistent with under-determination of the *fit*, rather than a limitation of the atlas representation itself. A direct conditioning analysis remains necessary to confirm this mechanism.

Table 11. Effect of raising the chart degree across all four tessellations: ratio of the degree-5 to the degree-3 Gaussian-curvature MAE at 5000 vertices (a value above one means the higher degree is worse). The oracle improves or is unchanged on every tessellation (≤ 1.01) and the purely positional fits degrade everywhere; the surface attractor with the analytic target sits between the two, behaving like the oracle except on the regular torus. Produced by `run_dof.py` (`Exp_dof/dof_table.csv`).

tessellation	oracle	strong	weak	sa (PN)	sa (torus)
regular	**0.92**	3.71	3.71	4.37	2.98
anisotropic	1.01	4.00	4.00	2.12	**0.86**
irregular	**0.85**	2.61	2.61	1.78	1.63
irregular anisotr.	**0.88**	4.19	4.19	1.86	0.98

8.5 Discussion of the Functional Maps experiment

The functional-maps experiment (Section 7.6) is a deliberate contrast to the curvature study. There the atlas is asked for a *derivative* of the surface; here it is asked only to *interpolate* a scalar field between the vertices. The two requests probe opposite regimes of the same construction, and the outcomes bracket a single message.

In this experiment, evaluating the signals through the smooth atlas yields a functional map numerically almost indistinguishable from the one obtained by plain linear interpolation — the two agree to the fourth decimal on every metric — and neither improves on the vertex-based discrete baseline. The reason is mundane and worth stating plainly: the mesh triangles are small, and over a small triangle the linear interpolant of a smooth, low-frequency signal is already very close to its smooth evaluation. The functional map is built from the first few Laplace–Beltrami eigenfunctions and from heat-kernel descriptors, all of them slowly varying; there is simply little sub-triangle structure for a smoother reconstruction to recover at the tested resolution and truncation, so the recovered map changes only in the fourth decimal.

This contrasts with the curvature experiments. Curvature is a *second derivative* and is sensitive to sub-triangle behaviour that linear interpolation cannot represent; the spectral signals used here are low-frequency order-zero data for which piecewise-linear interpolation is already accurate at the sampled scale. The combined evidence therefore suggests that smooth reconstruction is more likely to justify its cost for differential quantities than for interpolation of already well-resolved low-frequency signals. Testing other resolutions and signal families is necessary before generalising this conclusion. The directions these findings open — damping the artefact at its source, completing the convergence analysis, and asking systematically where a smooth atlas earns its cost — are collected in Chapter 10.

PART III

Conclusions

9. Conclusions

This thesis asked whether a C^k atlas fitted to a triangle mesh from vertex positions alone — order-zero information — also recovers a second-order differential quantity, the Gaussian curvature. The answer, measured on a controlled testbed rather than argued, has three parts.

Selected atlas reconstructions remain robust where the discrete error does not improve. Across the three tested resolutions, the strong, weak, and analytically targeted reconstructions retain decreasing vertex-curvature errors on the two irregular tessellations, while the angle-defect error does not decrease and can grow. They finish an order of magnitude below the discrete baseline. This is the main positive empirical result; it is not universal across all atlas configurations. For example, the unregularised strong and weak fits stall on the finest regular mesh, and the PN-targeted surface attractor degrades on irregular inputs. On well-shaped meshes the discrete estimator remains excellent and the atlas brings no advantage, an honest negative control. Section 8.1 therefore treats convergence conservatively. Xu's result concerns a uniformly well-posed *local quadratic* least-squares fit, not DALa's degree-three global variational solve. The common local viewpoint makes an analogous mechanism plausible at the vertices, while the different geometry proxies and coefficient estimators leave the required stability and approximation bounds open. The chapter consequently offers a conjecture and a proof sketch, not a convergence theorem for DALa.

Positional fidelity does not imply differential fidelity. This is the negative result, and the reason the question was worth asking. Sampling inside the triangles rather than at the vertices shows a median positional error of order 10^{-5} — on a torus of tube radius 0.4, a relative error below 0.01% — accompanied by a curvature error three orders of magnitude larger, the two being correlated only loosely. A reconstruction can match the sampled positions very closely and still misrepresent how it bends. The same asymmetry appears in the surface-attractor experiment: the PN-triangle target, built from vertex normals, degrades the reconstructed *position* by an order of magnitude on irregular meshes and the curvature by a factor of thirty. Compared with the analytically targeted run, this points to the surrogate as the main source of the degradation.

Blending contributes to the gap and provides a target for regularisation. Writing the reconstruction in common triangle coordinates as a partition of unity of the pulled-back charts and differentiating twice separates the error into a weighted average of the charts' own second-order errors and two terms that couple the *derivatives of the weights* with the disagreement between adjacent charts. The observed patterns are consistent with that decomposition. The blending terms vanish on the entire vertex plateaus, where one weight is identically one and the others are identically zero. The interior samples instead localise the larger errors in the multi-chart transition regions, consistently with the terms involving weight derivatives. Across most tested refinements these bands shrink as the local charts become more consistent. Three controlled experiments further support the role of both factors in the product-rule decomposition.

Changing only how the gate is applied — to the blending weight, for which C^k compatibility is proved, or to the barycentric coordinate, which is softer but is not covered by that proof — moves the curvature by up to a factor of 2.7 while leaving the position and the two non-blended estimators untouched to four decimals. The penalty is largest on the regular torus, where the fit is most under-determined, and vanishes on the irregular one, where the chart disagreements dominate regardless. Refining the quadrature that assembles the fit removes an apparent error floor, with the gain concentrated in a single step. This pattern is consistent with a sampling threshold for the narrow transition bands; their physical width depends on both the parametric gate and the triangle Jacobian. The gate and quadrature experiments therefore support the same numerical interpretation without establishing a universal cost of compatibility. Finally, regularising the fit so that neighbouring charts agree in value and first jet at the vertices attacks the other factor directly, improves every variational estimator at fine resolution by up to a factor of ten, removes the reversal of the unregularised regular-torus sequence, and does so at a sixteenth of the integration cost of the equivalent quadrature refinement. Damping the disagreement is cheaper than resolving the weights.

A fourth experiment probes the residual limitation. Raising the polynomial degree of the charts degrades every purely positional fit while leaving the oracle unchanged or better: the extra coefficients pay off when a correct target constrains them and are decided by the energy rather than by the data when it does not. The result points to under-determination of the *fit* by its effective quadrature constraints, rather than to the atlas representation itself; a direct conditioning analysis remains open.

Smoothness is not uniformly valuable. The last experiment carries the atlas into functional maps, where the same spectral signals are evaluated at the same sample points through linear interpolation and through the smooth reconstruction. The two agree to the fourth decimal on every metric, and neither improves on the plain vertex-based baseline, at a cost three orders of magnitude higher for the atlas. The reason is mundane and worth stating: the mesh triangles are small and the signals — the first Laplace–Beltrami eigenfunctions and heat-kernel descriptors — are slowly varying, so there is little sub-triangle structure for a smoother reconstruction to recover. Read together with the curvature results, this suggests a practical distinction: smooth reconstruction is more likely to justify its cost for differential quantities than for interpolation of already well-resolved, low-frequency signals. Other resolutions and signal families must be tested before this observation can be generalised.

On the scope of these conclusions. The C^k compatibility of the DALa parametrisation and its single-chart vertex plateaus are proved properties of the construction. The product-rule decomposition of the blended error is also exact, but the estimates needed to turn it into a convergence statement have not been proved for the global variational solve or for its numerical quadrature. The convergence language used here therefore describes finite experimental sequences, not an established asymptotic theorem. The testbed is a single analytic surface, chosen because its curvature is known exactly and its tessellations can be controlled; extending the account to other surfaces, other differential quantities and higher orders is precisely what Chapter 10 proposes. The product-rule decomposition makes explicit that a partition of unity, nearly invisible at order zero, can still influence the second-order behaviour through its derivatives.

10. Future Works

The experiments of Part II leave the atlas at a precise point: a positional reconstruction that estimates a second-order quantity accurately in selected well-resolved configurations, that exhibits a localised curvature artefact consistent with the blending weights, and that — on a task built from already-smooth, densely-sampled signals — gives no measurable improvement over piecewise-linear interpolation. Each of these outcomes suggests a continuation. The nearest are technical refinements of the reconstruction itself; the broader ones ask, systematically, *where* a smooth atlas earns its cost. The author intends to pursue the latter in doctoral work.

10.1 Damping the weight-derivative artefact

The analysis of Section 8.2 suggests a remedy. The additional blending terms couple weight derivatives, fixed once the gate is chosen, with the chart disagreements e_i and ∇e_i. The tested softer gate reduces the first factor but is not covered by the same compatibility proof. An alternative is to act on the other factor by adding a term to the fitting energy that penalises the charts for disagreeing — keeping e_i and ∇e_i small precisely where the weight derivatives are large, so that terms (b) and (c) of the error decomposition are damped at the source.

Section 8.3 carried out the first half of this programme: a regulariser penalising the mismatch of value and of the first jet at the vertices damps e_i and ∇e_i where it matters, improves every variational estimator at fine resolution — by up to $10\times$ — and restores a monotone decrease across the tested regular-torus refinements. What it does not yet do is act on the *second* jet, and it is the second derivatives of the chart disagreements that enter term (a) directly. Extending the penalty one order up is a small change given how the atlas exposes its coefficients, and the natural next experiment. Two questions follow it. The first is whether the penalty weights can be made adaptive — concentrated where $|\nabla w_i|$ is large, rather than uniform over the mesh — which the error decomposition suggests would buy the same accuracy for a smaller distortion of the fit. The second is why the regulariser *hurts* on the coarsest anisotropic meshes (Table 9): the counting argument of Section 8.4 suggests an answer, but it has not been tested.

10.2 Completing the convergence picture

The positive empirical result of Part II is that selected DALa reconstructions retain decreasing vertex-curvature errors on the tested irregular refinements, where the angle-defect error does not decrease. Section 8.1 does not present this as a theorem for DALa. It formulates a conjectural connection with Xu's local polynomial fitting and identifies what would have to be proved before the experimental trend could be given an asymptotic interpretation.

The first step is to extend the local comparison beyond Xu's specific setting. Xu estimates quadratic coefficients by least squares from the piecewise-linear geometry. DALa may instead use a PN-triangle or analytic geometry proxy, or may determine the coefficients through a global variational problem with no separate proxy beyond its positional constraints. A convergence analysis should establish which of these coefficient estimators are stable and consistent on the rescaled vertex domains. Replacing the PN target by samples of the piecewise-linear input would provide a useful intermediate case close to Xu; the harder step is to control how the variational energies, global coupling, regularisation, and quadrature perturb that local fit.

The second step is to investigate whether such a result can be made independent of a fixed polynomial degree and extended to differential operators of arbitrary order. Xu's analysis uses quadrics and consequently addresses quantities only up to second order. Since the triangular Bernstein basis of degree d spans the complete space $\mathcal{P}_d$, DALa offers a natural setting in which to ask how the approximation order, differential order, conditioning, and number of samples should scale together as d increases. The degree experiment of Section 8.4 warns that a larger approximation space is useful only when the geometry supplies enough effective constraints.

A final transfer step would connect pointwise vertex estimates to the blended surface. It must retain the dependence of the weight derivatives on the gate, transition maps, and triangle Jacobians; compare the quadrature-assembled and exact variational minimisers; and impose uniform immersion bounds before passing from derivative errors to curvature. Testing the resulting predictions on other analytic surfaces and other differential quantities would separate atlas error from the error introduced by the consuming algorithm.

10.3 Hybrid discrete–smooth algorithms

The two halves of Part II — curvature, where selected smooth reconstructions improve the estimate, and functional maps, where smooth signal evaluation gives no measurable benefit (Section 8.5) — bracket the question that a broader programme should answer directly: *where* does the continuous, off-vertex information of the atlas improve accuracy or robustness, and where does it merely add cost? The results here suggest that the answer is neither “always” nor “never” but a *hybrid*: keep the mature discrete algorithms, which are efficient and often already accurate enough, and let the atlas intervene only for the geometric quantities that genuinely benefit from being known between the vertices.

Two applications make good first tests. In *remeshing*, the discrete pipeline is left unchanged while the geometric quantities that steer it — curvatures, gradients, sizing fields — are evaluated on the atlas; comparing the output at a fixed pipeline isolates the contribution of the smooth information. In *geodesic distance*, a robust discrete solver supplies the distance field, the atlas reconstructs it smoothly, and its gradients and higher derivatives are then read off at arbitrary points — a concrete instance of the discrete solver and the smooth model playing to their respective strengths. Both are expected to matter most in the *low-resolution* regime, where the mesh samples a quantity too sparsely and knowing it inside each triangle is worth the reconstruction. The closed-form Laplacian of the atlas belongs here too: whether its eigenfunctions, computed directly rather than reconstructed from a discrete operator, change the outcome of the functional-maps experiment is a question this thesis sets up but does not settle.

10.4 A reproducible benchmark

Underlying all of the above is the need for a controlled way to compare discrete, smooth and hybrid strategies across accuracy, robustness and cost. The testbed of this thesis — analytic surfaces tessellated in several qualitatively different ways, with errors measured against a closed-form ground truth — is a first step in that direction. Growing it into a reproducible benchmark over several families of meshes and several differential quantities would turn isolated observations into validated results, and would let the choice between a discrete estimator, a smooth reconstruction and a hybrid of the two rest on measured criteria rather than on intuition.

PART IV

Appendices

A. Elements of topology

This appendix recalls all the topological notions used throughout the thesis, with the aim of making the exposition self-contained. The material is standard and it's the language in which the definitions of surface and manifold (Chapter 1) are written, and they also fix what it means for two shapes to be «the same» before one looks for a map between them.

The point-set material follows [20], the abstract-manifold notions (Hausdorff spaces, second-countable spaces, partitions of unity) follow [21], while the surface invariants (genus, Euler characteristic, Gauss–Bonnet theorem) are treated in [2]. The inverse function theorem is proved here in the classical form of [26]; [24, 25] cover the same ground, together with the linear algebra recalled in Appendix B.

A.1 Open sets, maps and spaces

Definition A.1 (Open and closed sets, neighbourhoods)**.** In $\mathbb{R}^n$ a set A is *open* if for every $p \in A$ there is a ball $B(p,\varepsilon) = \{x : \|x - p\| < \varepsilon\} \subseteq A$. A *closed* set is the complement of an open one, and a *neighbourhood* of p is an open set containing p.

Definition A.2 (Subspace (induced) topology)**.** Let $S \subseteq \mathbb{R}^n$. The *subspace topology* on S declares $A \subseteq S$ open if $A = S \cap U$ for some U open in $\mathbb{R}^n$. This is the topology that turns a regular surface $S \subseteq \mathbb{R}^3$ into a topological space in its own right and makes each parametrization a homeomorphism onto an open subset of S; without it the requirements in the definition of a regular surface would be meaningless.

Definition A.3 (Continuous map)**.** A map $f : A \subseteq \mathbb{R}^n \to \mathbb{R}^m$ is *continuous* if the preimage $f^{-1}(V)$ of every open set $V \subset \mathbb{R}^m$ is open in A.

Definition A.4 (Homeomorphism)**.** Let X, Y be topological spaces. A map $f : X \to Y$ is a *homeomorphism* if it is bijective, continuous, and has a continuous inverse f^{-1}. Two spaces related by a homeomorphism are *topologically identical* (they have the same open sets). It is the topological C^0 analogue of what a diffeomorphism is for the class C^∞.

Definition A.5 (Diffeomorphism)**.** Let $U, V \subset \mathbb{R}^n$ be open sets. A map $f : U \to V$ is a *diffeomorphism of class* C^k if it is bijective, of class C^k, with C^k inverse.

The following theorem is the reason a *local* diffeomorphism can be produced from a purely infinitesimal datum (invertibility of a single differential). It is used twice in this thesis: in the proof of Proposition 1.2 (where it upgrades a continuous inverse to a smooth one), and implicitly whenever a chart is inverted.

Theorem A.6 (Inverse function theorem)**.** Let $U \subseteq \mathbb{R}^n$ be open, $f : U \to \mathbb{R}^n$ of class C^k with $k \geq 1$, and let $p \in U$ be such that the differential $\mathrm{d}f_p$ is invertible. Then there exist open neighbourhoods $V \ni p$ and $W \ni f(p)$ such that $f|_V : V \to W$ is a diffeomorphism of class C^k, and

$$\mathrm{d}\big(f^{-1}\big)_{f(p)} \;=\; \big(\mathrm{d}f_p\big)^{-1}.$$

Proof. The proof is divided into six steps.

Reduction. Replacing f by $x \mapsto (\mathrm{d}f_p)^{-1}\big(f(x+p) - f(p)\big)$ changes neither the hypotheses nor the conclusion, since composing with an invertible linear map and a translation preserves the C^k class and maps diffeomorphisms to diffeomorphisms. We may therefore assume $p = 0$, $f(0) = 0$ and $\mathrm{d}f_0 = \mathbf{I}$.

Step 1 — Injectivity near 0. Put $g(x) = x - f(x)$, so that $\mathrm{d}g_0 = 0$. Since f is C^1, the map $x \mapsto \mathrm{d}g_x$ is continuous, so there is $r > 0$ with $\|\mathrm{d}g_x\| \leq \frac{1}{2}$ for all $x \in \bar{B}(0, r) \subseteq U$. The mean value inequality on the convex set $\bar{B}(0, r)$ gives

$$\|g(x) - g(y)\| \;\leq\; \tfrac{1}{2}\,\|x - y\|, \tag{A.1}$$

and therefore, writing $f = \mathrm{id} - g$,

$$\|f(x) - f(y)\| \;\geq\; \|x - y\| - \|g(x) - g(y)\| \;\geq\; \tfrac{1}{2}\,\|x - y\|. \tag{A.2}$$

In particular f is injective on $\bar{B}(0, r)$.

Step 2 — Surjectivity onto a ball. Let $W = B(0, r/2)$ and fix $y \in W$. Solving $f(x) = y$ is the same as finding a fixed point of

$$\varphi_y(x) \;=\; y + g(x) \;=\; y + x - f(x).$$

The map φ_y sends $\bar{B}(0, r)$ into itself: using $g(0) = 0$ and (A.1),

$$\|\varphi_y(x)\| \leq \|y\| + \|g(x) - g(0)\| \leq \tfrac{r}{2} + \tfrac{1}{2}\|x\| \leq r.$$

By (A.1) it is a $\frac{1}{2}$-contraction, and $\bar{B}(0, r)$ is a complete metric space, so Banach's fixed-point theorem gives a unique $x \in \bar{B}(0, r)$ with $f(x) = y$.

Setting $V = f^{-1}(W) \cap B(0, r)$, which is open because f is continuous, the restriction $f : V \to W$ is a bijection.

Step 3 — Continuity of the inverse. Inequality (A.2) reads, in terms of $u = f(x)$ and $v = f(y)$,

$$\|f^{-1}(u) - f^{-1}(v)\| \;\le\; 2\,\|u - v\|,$$

so f^{-1} is Lipschitz, hence continuous.

Step 4 — Differentiability of the inverse. Shrinking r if necessary we may assume $\mathrm{d}f_x$ invertible for every $x \in V$: the invertible matrices form an open set and $x \mapsto \mathrm{d}f_x$ is continuous. Fix $y_0 = f(x_0)$ with $x_0 \in V$, put $A = \mathrm{d}f_{x_0}$, and let $y \in W$, $x = f^{-1}(y)$. Then

$$f^{-1}(y) - f^{-1}(y_0) - A^{-1}(y - y_0) \;=\; -A^{-1}\Big[\, f(x) - f(x_0) - A\,(x - x_0) \Big].$$

The bracket is $o(\|x - x_0\|)$ because f is differentiable at x_0, and $\|x - x_0\| \le 2\|y - y_0\|$ by Step 3; hence the left-hand side is $o(\|y - y_0\|)$. This is exactly the statement that f^{-1} is differentiable at y_0 with $\mathrm{d}(f^{-1})_{y_0} = \big(\mathrm{d}f_{x_0}\big)^{-1}$.

Step 5 — Smoothness of the inverse. By Step 4,

$$\mathrm{d}\big(f^{-1}\big)_y \;=\; \Big(\mathrm{d}f_{f^{-1}(y)}\Big)^{-1},$$

a composition of three maps: f^{-1}, which is continuous by Step 3; $x \mapsto \mathrm{d}f_x$, which is C^{k-1} by hypothesis; and matrix inversion, which is smooth on the invertible matrices, its entries being rational functions of the entries with non-vanishing denominator by Cramer's rule. Hence $\mathrm{d}(f^{-1})$ is continuous, i.e. f^{-1} is C^1. Assuming inductively that f^{-1} is C^m for some $m < k$, the same identity exhibits $\mathrm{d}(f^{-1})$ as C^m, so f^{-1} is C^{m+1} and after k steps, f^{-1} is C^k. $\square$

Observation A.7 (Abstract topological space)**.** For manifolds one forgets the distance and keeps only the notion of open set. A *topological space* is a set X with a family of subsets (the open sets) closed under arbitrary unions and finite intersections, with $\varnothing$ and X open; continuity and homeomorphism are then defined as above, through preimages. Two technical hypotheses rule out pathologies: the space is *Hausdorff* (distinct points have disjoint neighbourhoods, so limits are unique) and *second countable* (the topology is generated by a countable family of open sets). Both hold for $\mathbb{R}^n$ and for every surface in $\mathbb{R}^3$.

A.2 Covers, compactness and surface invariants

Definition A.8 (Open cover, compactness, connectedness)**.** An *open cover* of X is a family of open sets whose union is X. A set $C \subseteq \mathbb{R}^n$ is *compact* if it is closed and bounded (equivalently, every open cover admits a finite subcover), and *connected* if it cannot be split into two disjoint non-empty open sets.

Observation A.9 (Partition of unity)**.** On a second-countable Hausdorff manifold, subordinate to any open cover $\{U_i\}$ there exists a *partition of unity*: smooth functions $\{\rho_i\}$ with $0 \leq \rho_i \leq 1$, each supported inside one U_i, and $\sum_i \rho_i \equiv 1$. This is the device that glues per-chart (local) constructions into a single global one, and is the continuous archetype of the local-to-global blending used on meshes.

The existence of such functions is not obvious — a non-zero analytic function cannot vanish on an open set, so the ρ_i must be smooth but *not* analytic. The following construction, classical and elementary [21, 25], produces them, and it is worth recording here because the blending weights of Chapter 4 are built on exactly this pattern: a smooth function that is identically one on a core, identically zero outside a slightly larger set, and monotone in between.

Lemma A.10 (Smooth bump function)**.** Let $0 < a < b$. There exists $\chi \in C^\infty(\mathbb{R}^n)$ with $0 \leq \chi \leq 1$, $\chi \equiv 1$ on the closed ball $\bar{B}(0, a)$ and $\chi \equiv 0$ outside $B(0, b)$.

Proof. Start from the one-variable function

$$f(t) = \begin{cases} e^{-1/t}, & t > 0, \\ 0, & t \leq 0, \end{cases}$$

which is smooth on all of $\mathbb{R}$: every derivative of $e^{-1/t}$ is of the form $P(1/t)\, e^{-1/t}$ for a polynomial P, and $e^{-1/t} \to 0$ faster than any power of $1/t$ diverges as $t \downarrow 0$, so all derivatives extend continuously by 0 at the origin. Note that f vanishes on a half-line while being non-zero elsewhere: this is precisely what an analytic function cannot do, and it is the whole point of the construction. From f build

$$g(t) = \frac{f(b-t)}{f(b-t) + f(t-a)}$$

and see that the denominator never vanishes, since for any t at least one of $b - t$ and $t - a$ is positive (as $a < b$), so g is smooth; moreover $g \equiv 1$ for $t \leq a$, $g \equiv 0$ for $t \geq b$, and $0 \leq g \leq 1$ throughout. Setting $\chi(x) = g(\|x\|)$ gives the claim: χ is smooth away from the origin as a composition, and near the origin it is constantly 1, so the non-smoothness of the norm at 0 is harmless. $\square$

Observation A.11 (From bumps to a partition of unity)**.** Given the cover $\{U_i\}$, one covers the manifold by coordinate balls each contained in some U_i, extracts a countable, locally finite subcover using second countability, and places a bump χ_i of Lemma A.10 on each. Local finiteness makes $\sigma = \sum_i \chi_i$ a well-defined smooth function, positive because the balls cover the manifold, and the normalisation

$$\rho_i \;=\; \frac{\chi_i}{\sigma}$$

gives the partition of unity. Two features of this last step reappear verbatim in the reconstruction studied in this thesis: the weights are obtained by *normalising* a family of bumps, so they sum to one by construction, and each is supported where its bump is. The price is that ρ_i inherits the derivatives of χ_i divided by σ — the transition region is where the weights vary, and the narrower it is, the larger those derivatives become. That elementary trade-off is what Chapter 8 measures on the curvature of the blended surface.

Definition A.12 (Topological invariant)**.** A quantity attached to a space is a *topological invariant* if homeomorphic spaces share it. Connectedness, compactness, and the Euler characteristic below are invariants; an invariant that differs already certifies that two shapes are *not* homeomorphic.

Definition A.13 (Closed surface, genus, Euler characteristic)**.** A surface that is compact and has no boundary is called *closed.* A closed orientable surface is determined, up to homeomorphism, by a single integer, its *genus* g (the number of handles: $g = 0$ for the sphere, $g = 1$ for the torus). For a triangulation with V vertices, E edges and F faces the *Euler characteristic*

$$\chi = V - E + F = 2 - 2g \tag{A.3}$$

is a topological invariant, independent of the chosen triangulation.

sphere: $g = 0,\ \chi = 2$ torus: $g = 1,\ \chi = 0$

Figure 26. A closed orientable surface is classified, up to homeomorphism, by its genus g (number of handles); the Euler characteristic $\chi = 2 - 2g$ is a topological invariant. The shapes matched in the experiments are genus-0 surfaces.

Observation A.14 (Why topology precedes matching)**.** The Gauss–Bonnet theorem ties this invariant to curvature, $\int_S K \, \mathrm{d}A = 2\pi\chi$, linking geometry and topology on a closed surface. It also explains a standing assumption of the matching pipeline: the shapes compared through a functional map (for instance the human meshes of the FAUST dataset) are closed, orientable surfaces of the *same* genus. A near-isometry, and hence a meaningful correspondence, can only be sought between shapes that are topologically compatible to begin with; an invariant mismatch would make the very notion of a low-distortion map ill posed.

B. Matrices properties

This appendix collects the matrix-algebra facts invoked, often implicitly, in the functional-maps construction and in the discrete operators of the preliminary chapters. It is a quick reference only: the operators themselves (Laplace–Beltrami, mass and stiffness matrices, the functional map) are defined where they are used and are not repeated here.

We follow [22] for the symmetric-matrix theory (spectral theorem, positive definiteness) and [23] for the generalized eigenproblem, the Cholesky factorisation, the pseudoinverse and least squares.

We write $\mathbf{A} \in \mathbb{R}^{m\times n}$ for a real matrix, $\mathbf{A}^\top$ for its transpose, $\mathbf{I}$ for the identity, and $\langle x, y\rangle = x^\top y$ for the Euclidean inner product on $\mathbb{R}^n$.

B.1 Some classical results

Definition B.1 (Symmetric matrix)**.** A square matrix $\mathbf{A} \in \mathbb{R}^{n\times n}$ is *symmetric* if

$$\mathbf{A}^\top = \mathbf{A}$$

and it is the matrix counterpart of a self-adjoint operator:

$$\langle \mathbf{A}x, y\rangle = \langle x, \mathbf{A}y\rangle, \quad \forall x, y.$$

Definition B.2 (Orthogonal matrix)**.** A matrix $\mathbf{Q} \in \mathbb{R}^{n\times n}$ is *orthogonal* if

$$\mathbf{Q}^\top\mathbf{Q} = \mathbf{Q}\mathbf{Q}^\top = \mathbf{I},$$

equivalently if its columns form an orthonormal basis. Orthogonal matrices preserve the inner product, $\langle \mathbf{Q}x, \mathbf{Q}y\rangle = \langle x, y\rangle$, hence lengths and angles; their inverse is their transpose.

Definition B.3 (Eigenvalues and eigenvectors)**.** A scalar λ and a non-zero vector v form an *eigenpair* of $\mathbf{A}$ if $\mathbf{A}v = \lambda v$. The set of eigenvalues is the *spectrum* of $\mathbf{A}$.

Theorem B.4 (Spectral theorem for symmetric matrices)**.** Every symmetric $\mathbf{A} \in \mathbb{R}^{n\times n}$ has real eigenvalues and an orthonormal basis of eigenvectors: there exist an orthogonal $\mathbf{Q}$ and a diagonal $\mathbf{D} = \operatorname{diag}(\lambda_1, \dots, \lambda_n)$ with

$$\mathbf{A} = \mathbf{Q}\,\mathbf{D}\,\mathbf{Q}^\top = \sum_{i=1}^{n} \lambda_i \, q_i q_i^\top. \tag{B.1}$$

This is the algebraic counterpart of the eigen-decomposition of the Laplace–Beltrami operator: a symmetric operator admits a real spectrum and an orthonormal eigenbasis, which is what makes the spectral embedding well posed.

Proof. The proof is divided into two steps.

Step 1 — Reality of the spectrum. Let $\mathbf{A}v = \lambda v$ with $v \neq 0$, allowing complex entries for the moment. Taking the conjugate transpose, $\bar{v}^\top \mathbf{A} v = \lambda\, \bar{v}^\top v$, while the same quantity equals $\overline{\lambda}\, \bar{v}^\top v$ because $\mathbf{A}$ is real and symmetric, hence self-adjoint. Since $\bar{v}^\top v = \|v\|^2 > 0$, we get $\lambda = \overline{\lambda}$, so $\lambda \in \mathbb{R}$; the corresponding eigenvector can then be chosen real.

Step 2 — Existence of an orthonormal eigenbasis (by induction on n). For $n = 1$ there is nothing to prove. For $n > 1$, the Rayleigh quotient $v \mapsto v^\top \mathbf{A} v$ is continuous on the unit sphere, which is compact, so it attains a maximum at some q_1 with $\|q_1\| = 1$. A Lagrange-multiplier argument at that maximum gives $\mathbf{A}q_1 = \lambda_1 q_1$ with $\lambda_1 = q_1^\top \mathbf{A} q_1$: the maximiser is an eigenvector.
Consider now $H = q_1^\perp$, of dimension $n-1$. It is $\mathbf{A}$-*invariant*: if $w \perp q_1$ then

$$\langle \mathbf{A}w, q_1\rangle = \langle w, \mathbf{A}q_1\rangle = \lambda_1 \langle w, q_1\rangle = 0,$$

where the first equality is symmetry — this is the step that fails for a non-symmetric matrix, and the reason the theorem is about symmetric ones. The restriction $\mathbf{A}|_H$ is again symmetric on an $(n-1)$-dimensional space, so by the inductive hypothesis it admits an orthonormal eigenbasis $q_2, \dots, q_n$ of H. Together with q_1 these form an orthonormal basis of $\mathbb{R}^n$ made of eigenvectors; collecting them as the columns of $\mathbf{Q}$ gives $\mathbf{Q}^\top\mathbf{Q} = \mathbf{I}$ and $\mathbf{AQ} = \mathbf{QD}$, which is the stated factorisation. □

Definition B.5 (Positive (semi)definiteness)**.** A symmetric matrix $\mathbf{A}$ is *positive semidefinite* (PSD) if $x^\top \mathbf{A} x \geq 0$ for all x, and *positive definite* (SPD) if the inequality is strict for $x \neq 0$. Equivalently, all eigenvalues are ≥ 0 (resp. > 0).

Property B.6 (Energy and inner product)**.** If $\mathbf{B}$ is SPD then $\langle x, y\rangle_{\mathbf{B}} := x^\top \mathbf{B} y$ is an inner product. The mass matrix plays exactly this role: it turns the vector of nodal values of a function into its L^2 inner product on the mesh, so that "orthonormal" descriptors are orthonormal *with respect to* $\mathbf{B}$.

Definition B.7 (Generalized eigenproblem)**.** Given a symmetric $\mathbf{A}$ and an SPD $\mathbf{B}$, the *generalized eigenproblem* seeks pairs (λ, v) with

$$\mathbf{A}v = \lambda\,\mathbf{B}v. \tag{B.2}$$

Property B.8 ($\mathbf{B}$-orthogonal eigenbasis)**.** The generalized problem above has real eigenvalues and a basis of eigenvectors that is $\mathbf{B}$-*orthonormal*, i.e. $v_i^\top \mathbf{B} v_j = \delta_{ij}$. Writing $\mathbf{B} = \mathbf{L}\mathbf{L}^\top$ (Cholesky), the substitution $w = \mathbf{L}^\top v$ reduces it to the ordinary symmetric problem $\mathbf{L}^{-1}\mathbf{A}\mathbf{L}^{-\top} w = \lambda w$, so the spectral theorem applies. This is the form solved when diagonalizing the stiffness matrix against the mass matrix.

B.2 Pseudoinverse and least squares

Definition B.9 (Moore–Penrose pseudoinverse)**.** For $\mathbf{A} \in \mathbb{R}^{m\times n}$ the *pseudoinverse* $\mathbf{A}^+$ is the unique matrix satisfying the four *Penrose conditions*:

1. $\mathbf{A}\mathbf{A}^+\mathbf{A} = \mathbf{A}$ (weak inverse),
2. $\mathbf{A}^+\mathbf{A}\mathbf{A}^+ = \mathbf{A}^+$ (reflexivity),
3. $(\mathbf{A}\mathbf{A}^+)^\top = \mathbf{A}\mathbf{A}^+$ ($\mathbf{A}\mathbf{A}^+$ is symmetric),
4. $(\mathbf{A}^+\mathbf{A})^\top = \mathbf{A}^+\mathbf{A}$ ($\mathbf{A}^+\mathbf{A}$ is symmetric).

When $\mathbf{A}$ has full column rank, these conditions are jointly satisfied by

$$\mathbf{A}^+ = (\mathbf{A}^\top\mathbf{A})^{-1}\mathbf{A}^\top.$$

Property B.10 (Least-squares solution)**.** The minimizer of $\|\mathbf{A}x - b\|_2$ is $x^\star = \mathbf{A}^+b$; for several right-hand sides $\mathbf{B}$, the minimizer of $\|\mathbf{A}\mathbf{X} - \mathbf{B}\|_F$ is $\mathbf{X}^\star = \mathbf{A}^+\mathbf{B}$. The functional map is estimated in exactly this way: it is the matrix $\mathbf{C}$ that best maps the source descriptor coefficients onto the target ones in the truncated spectral basis, a linear least-squares problem solved through the pseudoinverse.

Proof. Assume first that $\mathbf{A}$ has full column rank, the case relevant here, so that $\mathbf{A}^+ = (\mathbf{A}^\top\mathbf{A})^{-1}\mathbf{A}^\top$. The key is a geometric observation: as x ranges over $\mathbb{R}^n$, the vector $\mathbf{A}x$ ranges over the column space $\operatorname{im}\mathbf{A}$, so the problem is to find the point of that subspace closest to b — the orthogonal projection of b onto it. The minimiser is therefore characterised by requiring the residual to be orthogonal to the subspace,

$$\mathbf{A}x^\star - b \;\perp\; \operatorname{im}\mathbf{A} \qquad\iff\qquad \mathbf{A}^\top(\mathbf{A}x^\star - b) = 0,$$

which are the *normal equations* $\mathbf{A}^\top\mathbf{A}\,x^\star = \mathbf{A}^\top b$, with the unique solution $x^\star = \mathbf{A}^+b$. To confirm that this is a minimum, write any $x = x^\star + d$. Then

$$\|\mathbf{A}x - b\|_2^2 = \|(\mathbf{A}x^\star - b) + \mathbf{A}d\|_2^2 = \|\mathbf{A}x^\star - b\|_2^2 + \|\mathbf{A}d\|_2^2,$$

the cross term $2\,d^\top \mathbf{A}^\top(\mathbf{A}x^\star - b)$ vanishing by the normal equations. The right-hand side is minimal exactly when $\mathbf{A}d = 0$, i.e. $d = 0$ under the full-rank assumption; without it the minimiser is no longer unique, and $\mathbf{A}^+ b$ selects the one of smallest norm. For several right-hand sides the Frobenius norm decouples by columns,

$$\|\mathbf{A}\mathbf{X} - \mathbf{B}\|_F^2 = \sum_k \|\mathbf{A}x_k - b_k\|_2^2,$$

so each column is an independent instance of the problem and $\mathbf{X}^\star = \mathbf{A}^+\mathbf{B}$. $\square$

Bibliography

[1] M. P. do Carmo, *Riemannian Geometry*, Birkhäuser, 1992.

[2] M. P. do Carmo, *Differential Geometry of Curves and Surfaces (Revised & Updated Second Edition)*, Dover Publications Inc., 2016.

[3] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, L. Guibas, *Functional Maps: A Flexible Representation of Maps Between Shapes*, ACM Transactions on Graphics (Proc. SIGGRAPH) 31, 4, 2012.

[4] B. Lévy, *Laplace–Beltrami Eigenfunctions: Towards an Algorithm That Understands Geometry*, Proc. IEEE International Conference on Shape Modeling and Applications (SMI), 2006.

[5] R. M. Rustamov, *Laplace–Beltrami Eigenfunctions for Deformation Invariant Shape Representation*, Proc. Symposium on Geometry Processing (SGP), 225–233, 2007.

[6] J. Sun, M. Ovsjanikov, L. Guibas, *A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion*, Computer Graphics Forum (Proc. SGP) 28, 5, 1383–1392, 2009.

[7] P. Bérard, G. Besson, S. Gallot, *Embedding Riemannian Manifolds by Their Heat Kernel*, Geometric and Functional Analysis (GAFA) 4, 4, 373–398, 1994.

[8] T. Kato, *Perturbation Theory for Linear Operators*, Springer-Verlag, 1995.

[9] M. Ovsjanikov, Q. Mérigot, F. Mémoli, L. Guibas, *One Point Isometric Matching with the Heat Kernel*, Computer Graphics Forum (Proc. SGP) 29, 5, 1555–1564, 2010.

[10] M. Aubry, U. Schlickewei, D. Cremers, *The Wave Kernel Signature: A Quantum Mechanical Approach to Shape Analysis*, Proc. ICCV Workshops (4DMOD), 2011.

[11] M. Meyer, M. Desbrun, P. Schröder, A. H. Barr, *Discrete Differential-Geometry Operators for Triangulated 2-Manifolds*, in Visualization and Mathematics III, Springer, 35–57, 2003.

[12] U. Pinkall, K. Polthier, *Computing Discrete Minimal Surfaces and Their*

Conjugates, Experimental Mathematics 2, 1, 15–36, 1993.

[13] V. G. Kim, Y. Lipman, T. Funkhouser, *Blended Intrinsic Maps*, ACM Transactions on Graphics (Proc. SIGGRAPH) 30, 4, 2011.

[14] F. Bogo, J. Romero, M. Loper, M. J. Black, *FAUST: Dataset and Evaluation for 3D Mesh Registration*, Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014.

[15] A. Vlachos, J. Peters, C. Boyd, J. L. Mitchell, *Curved PN Triangles*, Proc. ACM Symposium on Interactive 3D Graphics (I3D), 2001.

[16] G. Xu, *Consistent approximations of several geometric differential operators and their convergence*, Applied Numerical Mathematics 69, 1–12, Elsevier, 2013, doi:10.1016/j.apnum.2013.02.002.

[17] F. Cazals, M. Pouget, *Estimating differential quantities using polynomial fitting of osculating jets*, Computer Aided Geometric Design 22 (2), 121–146, 2005.

[18] T. Djuren, U. Finnendahl, M. Kohlbrenner, M. Worchel, M. Alexa, *Interpolating Splines over Triangulated Surfaces by Blending Vertex-Centric Local Geometries*, Computers & Graphics, 2025.

[19] C. Mancinelli et al., *DALa: a Differentiable Atlas Language for Smooth Mesh Processing*, unpublished manuscript in preparation, 2026.

[20] J. R. Munkres, *Topology*, 2nd ed., Prentice Hall, 2000.

[21] J. M. Lee, *Introduction to Smooth Manifolds*, 2nd ed., Graduate Texts in Mathematics 218, Springer, 2013.

[22] R. A. Horn, C. R. Johnson, *Matrix Analysis*, 2nd ed., Cambridge University Press, 2012.

[23] G. H. Golub, C. F. Van Loan, *Matrix Computations*, 4th ed., Johns Hopkins University Press, 2013.

[24] E. Sernesi, *Geometria 1*, 2nd ed., Bollati Boringhieri, 2000.

[25] E. Sernesi, *Geometria 2*, Bollati Boringhieri, 1994.

[26] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed., McGraw–Hill, 1976.