

DISSERTATION

AN ENERGY-BASED FRAMEWORK FOR CONTACT-AWARE MESH DEFORMATION IN
CREATIVE SOFTWARE

Submitted by

Tristan Neighbors

Department of Mathematics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Summer 2026

Doctoral Committee:

Advisor: Clayton Shonkwiler

Alexander Hulpke
Wolfgang Bangerth
Sanjay Rajopadhye

Copyright by Tristan Neighbors 2026

All Rights Reserved

ABSTRACT

AN ENERGY-BASED FRAMEWORK FOR CONTACT-AWARE MESH DEFORMATION IN CREATIVE SOFTWARE

This thesis develops a customizable framework for a class of mesh deformation problems in creative software. In particular, it concerns problems where a mesh, here called the **conforming surface** or **conforming mesh**, is to be deformed in a way that

- (i) respects geometric contact or near-contact constraints defined by a fixed *target surface* and
- (ii) preserves the character of the mesh in a way that can be tuned by a set of interpretable parameters.

This fixed target surface is represented by a higher-order algebraic discrete signed distance function (ADSDF), while the conforming mesh carries Euclidean vertex positions together with auxiliary unit-quaternion rotational data. We refer to this process as **retargeting**. Retargeting is formulated as the minimization of an energy built from signed-distance evaluations and their derivatives, combining contact-aware energy terms with an energy term inspired by the **as-rigid-as-possible** (ARAP) technique [1] that preserves designer intent by promoting near-isometric deformations. Optional terms for normal alignment, conformity to prescribed geometric landmarks, smoothing, and related constraints can be added. The resulting optimization problem is posed intrinsically on a product space of Euclidean and manifold-valued variables.

The thesis makes three main contributions. First, it organizes mesh retargeting against signed-distance targets as a customizable energy optimization framework – the choice of energy terms determines the character of the final retargeted mesh. Second, it proves a local theorem for ADSDFs: on regular regions, least-squares (LSQ) polynomial samples inherit the same approximation order as Taylor polynomial samples, providing a mathematical justification for the practically im-

portant LSQ construction given by Valasek and Bán [2]. Third, it extends subdivision to manifold-valued data by replacing affine averages with Karcher means and proves continuous convergence by proximity to the underlying linear scheme.

The framework is illustrated through problem formulations for fitting a shirt around an aligned body, pose-retargeting a shirt with landmarked geometric features, and armature-driven contact resolution against a synthetic obstacle field. Together these examples show how higher-order signed-distance fields, customizable energies, and intrinsic optimization can be combined to preserve design intent while resolving geometric constraints.

ACKNOWLEDGEMENTS

I am thankful to my parents for their love and support throughout my graduate studies. They were always enthusiastic to hear about my work even if they didn't fully understand the details.

I am deeply grateful to my partner and wife-to-be, Kisa Kaster, for her warmth, patience, and encouragement, especially during times when stress got the better of me.

I am also grateful to my advisor, Clayton Shonkwiler, not only for his mentorship, but also for his encouragement during a period when I was on the verge of quitting. It seems unlikely that this work would have been completed without his patience and understanding.

Finally, I would like to thank my committee members, Alexander Hulpke, Wolfgang Bangerth, and Sanjay Rajopadhye, for their time, feedback, and thoughtful engagement with this thesis. Dr. Bangerth's expertise in numerical analysis was especially valuable, and this dissertation would be substantially poorer without his insights.

DEDICATION

I would like to dedicate this thesis to my companion and wife-to-be, Kisa Kaster.

TABLE OF CONTENTS

ABSTRACT	ii
ACKNOWLEDGEMENTS	iv
DEDICATION	v
LIST OF TABLES	ix
LIST OF FIGURES	x
Chapter 1 Introduction	1
1.1 The problems we are interested in	1
1.2 Problem formulation	2
1.3 The role of the ARAP term	5
1.4 Main results and organization of the thesis	6
Chapter 2 Retargeting Energies and SDFs	9
2.1 Signed-distance preliminaries	9
2.1.1 Signed distance functions, regular points, and the cut locus	10
2.1.2 The Hessian and curvature of level sets	11
2.2 The retargeting energy at a glance	12
2.3 SDF-driven terms	14
2.3.1 Distance, offset, and clearance	14
2.3.2 Normal alignment	15
2.4 Near-isometry and the ARAP term	16
2.4.1 Smooth rigidity and its rotational part	16
2.4.2 The discrete near-isometric energy	17
2.4.3 Relation to classical ARAP	18
2.5 Auxiliary terms	18
2.5.1 Landmarks	18
2.5.2 Auxiliary smoothing or regularization	19
2.5.3 Energy weights and regionwise compliance	19
2.6 What remains before optimization is possible	19
Chapter 3 Algebraic Discrete Signed Distance Fields	21
3.1 From DSDFs to ADSDFs	21
3.1.1 Discrete signed distance functions as a baseline	21
3.1.2 Higher-order sample fields and Taylor ADSDFs	22
3.1.3 Why the interpolants matter: derivatives and Hermite-type endpoint conditions	25
3.2 Representing, filtering, and constructing ADSDFs in practice	28
3.2.1 Local coordinates, global monomials, and why normalization matters	28
3.2.2 Cubical and simplicial filters	31
3.2.3 Least-squares ADSDFs from fine stencils	32
3.3 A local LSQ-to-Taylor theorem	35

3.3.1	Setup	35
3.3.2	Samplewise comparison	36
3.3.3	From samples to filtered fields	38
3.3.4	Limitations and bridge to the next chapter	41
Chapter 4	Intrinsic Optimization on Euclidean and Manifold-Valued Variables	42
4.1	Why manifold optimization appears in retargeting	42
4.2	Riemannian gradients, retractions, and descent	44
4.3	Lie groups with bi-invariant metrics	46
4.4	Working concretely in S^3	47
4.5	Quaternion identities for gradients and Jacobians	49
4.6	Gradients of the retargeting energies	52
4.6.1	Distance, clearance, and landmarks	52
4.6.2	Normal alignment	54
4.6.3	The ARAP term	54
4.6.4	Smoothing and Karcher-type terms	55
4.7	Karcher means as the recurring worked example	56
4.8	Why Gauss–Newton rather than Newton?	56
4.9	Riemannian Gauss–Newton and Levenberg–Marquardt	58
4.10	Residual formulation of the retargeting objective	59
4.11	What this chapter supplies for the rest of the thesis	61
Chapter 5	Subdivision and Multiresolution Refinement of Manifold-Valued Data	62
5.1	Why subdivision appears in this thesis	62
5.2	A visual introduction to subdivision surfaces	63
5.2.1	Repeated refinement and limit surfaces	63
5.2.2	Loop and Catmull–Clark	64
5.2.3	Regular and extraordinary vertices	65
5.3	Classical concepts we will use	66
5.3.1	Stationary affine-invariant schemes and local rules	66
5.3.2	Control rings, local subdivision matrices, and quotient spaces	66
5.3.3	Characteristic maps and imported smoothness facts	70
5.4	From affine averages to Karcher means	70
5.5	Proof roadmap	71
5.6	Continuous convergence by proximity	72
5.6.1	Hypotheses	73
5.6.2	Local expansion of the Karcher mean	74
5.6.3	One-step proximity	76
5.6.4	Contraction on regular patches	77
5.6.5	Extraordinary-ring contraction	78
5.6.6	Global decay and convergence of interpolants	79
5.7	Consequences and role in the thesis	82
Chapter 6	Illustrative Applications and Experimental Formulations	83
6.1	Common residuals and solver conventions	83

6.1.1	A smooth clearance residual	83
6.1.2	A weighted near-isometric term	84
6.1.3	Experimental protocol and reported quantities	85
6.2	Shirt fitting on an aligned body	86
6.2.1	Geometric setup and intended phenomenon	86
6.2.2	Optimization variables and energies	86
6.3	Pose-retargeted shirt with landmarked openings	88
6.3.1	Geometric setup and intended phenomenon	88
6.3.2	Landmark formulation	88
6.3.3	Optimization problem	90
6.4	Armature-driven tentacle in a synthetic obstacle field	91
6.4.1	Geometric setup and intended phenomenon	91
6.4.2	Armature-restricted search space	93
6.4.3	Energy and numerical problem	93
6.5	Robustness and computational cost	95
Chapter 7	Conclusion	97
7.1	Summary of contributions	97
7.2	Limitations	98
7.3	Future directions	98
7.3.1	Physically motivated dynamic extensions	98
7.3.2	Adaptive and multilevel manifold-valued subdivision	99
7.3.3	Adaptive and compressive higher-order SDFs	99
Appendix A	Reference Material on Smooth Signed Distance Functions	101
A.1	Basic facts	101
A.2	Characterization by Eikonal and closest-point properties	102
A.3	Exact formulas for common primitives	103
A.3.1	Sphere	103
A.3.2	Torus	104
A.3.3	Oriented plane	104
A.3.4	Point and segment primitives	105
A.4	Curvature of level sets from the Hessian	105
Appendix B	Quaternions	107
B.1	Algebra of quaternions	107
B.2	Geometry of quaternions	111
Bibliography		114

LIST OF TABLES

3.1	Coefficient counts for total-degree polynomial samples in three dimensions.	35
4.1	Derivative order required by Gauss–Newton versus full Newton for the most important residuals in this thesis. The derivative-count argument is one of the main reasons to favor Gauss–Newton in an ADSDF-based formulation.	57
6.1	Parameters for the aligned body/shirt clearance experiment.	87
6.2	Numerical behavior of the aligned body/shirt clearance experiment.	87
6.3	Parameters for the ADSDF body/shirt landmark retargeting experiment.	91
6.4	Numerical behavior of the ADSDF body/shirt landmark retargeting experiment. The initial row is measured after the harmonic displacement seed. The raw unseeded shirt has boundary landmark RMS error about 0.450 before retargeting.	92
6.5	Parameters for the ADSDF armature-driven tentacle experiment.	94
6.6	Numerical behavior of the ADSDF armature-driven tentacle experiment.	95

LIST OF FIGURES

1.1	(Left) A red shirt asset fitted to a T-posed body. (Right) Artist modifies body pose and the shirt is retargeted to the new pose.	2
1.2	Visualization of a vertex i and its dual $\star i$	5
2.1	One-ring neighborhood $\mathcal{N}_1(i)$ (yellow) of vertex i (teal).	13
3.1	One-dimensional Hermite-type blending functions.	28
3.2	(Blue) A coarse sample point. (Red) A fine stencil of additional samples around it. . .	33
4.1	A unit quaternion visualized by its action on a coordinate axis.	48
4.2	The Karcher mean of two unit quaternions q_1, q_2	51
5.1	(Left) A conforming surface is to be “shrinkwrapped” to a target monkey head surface. (Center) The procedure begins with a coarse resolution where the conforming surface cannot be cleanly fit to the target. (Right) The conforming surface is refined and is able to be cleanly shrinkwrapped to the target.	63
5.2	A mesh $\mathcal{M}^0$ is successively refined into $\mathcal{M}^1, \mathcal{M}^2, \dots$. In the limit, it becomes a smooth surface $\mathcal{M}^\infty$	64
5.3	Splitting a quadrilateral into four smaller quadrilaterals for Catmull–Clark subdivision. A point has been inserted at the center of the faces and at the midpoint of the 4 edges. .	65
5.4	Masks for the Catmull–Clark scheme.	67
5.5	(Blue) a particular vertex in a patch of a quadrilateral mesh. (Red) Its corresponding control ring under Catmull–Clark subdivision.	69
6.1	A shirt fitting example on the loaded body and shirt meshes. The initial shirt penetrates the body. The optimized result clears the body while preserving the overall shirt shape.	88
6.2	Diagnostics for the aligned body/shirt clearance experiment. The curves show normalized total energy, maximum clearance violation, and violating-vertex count over the active convergence window.	89
6.3	Landmark-guided shirt retargeting. (Left) A red shirt asset fitted to a T-posed body, with landmark features visualized in blue. (Right) The body pose and landmark positions are modified before optimization.	90
6.4	Diagnostics for the ADSDF body/shirt landmark retargeting experiment. The plotted curves show the seeded optimization after the harmonic landmark initialization. . . .	92
6.5	Armature-driven tentacle interaction with a synthetic obstacle field. The optimization removes the initial interpenetration by moving the low-dimensional control curve, rather than by freely editing every mesh vertex.	95
6.6	Diagnostics for the ADSDF armature-driven tentacle experiment. The maximum violation and violating-vertex count drop to zero early, after which the energy decreases as the armature settles within the remaining guide and preservation terms.	96

Chapter 1

Introduction

1.1 The problems we are interested in

A common task in 3D creative software is to deform one surface according to the geometry of another while respecting constraints that matter to a designer. We will call the fixed surface the **target surface** and the surface being deformed the **conforming surface** or **conforming mesh**. Sometimes the goal is simply to bring the conforming surface into contact with the target. More often a mesh asset, provided by a designer or artist, must be **retargeted** to a new surface while preserving seams, silhouette, texture layout, or some subjective aspect of design intent.

A garment-retargeting example captures the main difficulty. Suppose a shirt has been modeled for a base humanoid and must be adapted to a more muscular body. If the original shirt is placed on the new body in the appropriate pose without changing its geometry, parts of the garment may intersect the body mesh unrealistically—a phenomenon known in computer graphics as **clipping**. The task is then to deform the shirt so that it clears the body while preserving the visible structure of the asset: hems should remain plausible, sleeve cuffs should stay near the wrists, seams should not wander, and the resulting garment should still look like the same shirt.

The approach developed in this thesis combines two ingredients. First, the target surface is represented by a **signed distance function (SDF)** [3, Section 14.6]. Informally, an SDF is a scalar field whose zero set is the target surface and whose value at a point records the signed distance to that surface: positive outside and negative inside. Second, the conforming mesh is deformed by minimizing an energy built from the geometric information supplied by that field. This point of view is especially well suited to situations in which the target is fixed and the conforming surface is being edited, since the cost of building the SDF representation is then amortized over many evaluations.




Figure 1.1: (Left) A red shirt asset fitted to a T-posed body. (Right) Artist modifies body pose and the shirt is retargeted to the new pose.

The appeal of the SDF representation is computational. It supports fast geometric queries¹ and, in the higher-order algebraic constructions used here, it also provides derivatives analytically rather than through finite differences. This matters both mathematically and computationally. Surface normals and curvature information can be extracted directly from the chosen representation, and these evaluations are well suited to graphics hardware [2, 4, 5].

1.2 Problem formulation

We denote the target surface by $\mathcal{T}$ and the conforming surface by $\mathcal{C}$. The surface $\mathcal{C}$ is represented as a mesh. The surface $\mathcal{T}$ is converted to a discrete signed distance function $\hat{\Phi}_{\text{SDF}}$ using higher-order algebraic constructions [2]. A classical discrete signed distance function stores sampled distance values and interpolates them. An **algebraic discrete signed distance function (ADSDF)** stores local polynomial data instead. This is useful here because the gradients and Hessians of those local polynomials can be evaluated directly at query points, so the representation carries not only distances but also differential information about the target surface.

¹By a geometric query I mean the evaluation at a point in space of quantities such as signed distance to the target, a surface normal, or higher differential data.

The conforming mesh carries two kinds of per-vertex data. The first is Euclidean data, namely the vertex positions

$$\mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_N) \in (\mathbb{R}^3)^N.$$

The second is manifold-valued data

$$\mathbf{Q} = (\mathbf{q}_1, \dots, \mathbf{q}_N) \in \mathcal{M}^N.$$

Here $\mathcal{M}$ is a **Riemannian manifold**: a smooth manifold equipped with an inner product on each tangent space, varying smoothly from point to point. This gives a notion of gradient and geodesic motion that respects the geometry of the manifold itself. In the main application of this thesis we take

$$\mathcal{M} = S^3,$$

the manifold of unit quaternions. These quaternions encode per-vertex rotational information, which serves as the discrete analogue of the rotational part of a deformation.

The manifold S^3 is also a **Lie group**, that is, a smooth manifold with a smooth group multiplication law. This is useful computationally because quaternion multiplication and the quaternion exponential are explicit. Optimization updates can therefore be written directly in quaternion arithmetic rather than re-derived in local coordinates at every step [6–8].

The basic retargeting ingredients are a distance term and an ARAP-like near-isometry term. Depending on the application, one may also add terms for normal alignment, smoothing of the rotational field, adherence to landmarks, or nonpenetration barriers. By a **landmark** we mean a user- or algorithm-specified feature that should remain near a designated point or region, such as a

sleeve cuff near a wrist or the neckline of a shirt. A typical energy therefore has the form

$$\begin{aligned}\mathcal{E}(\mathbf{P}, \mathbf{Q}; \mathbf{T}) = & \underbrace{T_{\text{dist}}\mathcal{E}_{\text{dist}}(\mathbf{P}) + T_{\text{ARAP}}\mathcal{E}_{\text{ARAP}}(\mathbf{P}, \mathbf{Q})}_{\text{basic retargeting terms}} \\ & + \underbrace{T_{\text{norm}}\mathcal{E}_{\text{norm}}(\mathbf{P}, \mathbf{Q}) + T_{\text{smooth}}\mathcal{E}_{\text{smooth}}(\mathbf{P}, \mathbf{Q})}_{\text{alignment and regularity}} \\ & + \underbrace{T_{\text{land}}\mathcal{E}_{\text{land}}(\mathbf{P}) + T_{\text{barrier}}\mathcal{E}_{\text{barrier}}(\mathbf{P})}_{\text{application-dependent constraints}},\end{aligned}$$

where the energy weights $\mathbf{T}$ control the relative strength of the various terms. The notation $T_{\text{dist}}, T_{\text{ARAP}}, \dots$ is retained for continuity with the implementation, but these coefficients are only relative weights for residual blocks; no thermodynamic interpretation is intended. Different applications may omit some of these terms or add new ones without changing the basic optimization framework.

More generally, these energies are assembled from local evaluations of the target field and its derivatives:

$$\mathcal{E}(\mathbf{P}, \mathbf{Q}; \mathbf{T}) = \sum_{i \in \mathcal{V}} \hat{K}(\mathbf{P}, \mathbf{Q}, \hat{\Phi}_{\text{SDF}}(\mathbf{p}_i), \nabla \hat{\Phi}_{\text{SDF}}(\mathbf{p}_i), \nabla^2 \hat{\Phi}_{\text{SDF}}(\mathbf{p}_i), \dots) \text{area}(\star i).$$

Here $\text{area}(\star i)$ denotes a standard dual-area weight associated with vertex i , for example the area of a vertex-centered dual cell (figure 1.2). This expression makes the customizability of the framework explicit: once the target has been encoded as $\hat{\Phi}_{\text{SDF}}$ and the conforming mesh carries the rotational data $\mathbf{Q}$, one can change the collection of energies without changing the overall architecture of the method.

This formulation should be compared with the classical ARAP pipeline [1]. Standard ARAP methods typically alternate between updates of the vertex positions and updates of the local rotational data. In our method, the variables $\mathbf{P}$ and $\mathbf{Q}$ are updated within a single intrinsic descent step on the product space $(\mathbb{R}^3)^N \times \mathcal{M}^N$. This is the change that makes the optimization natu-




Figure 1.2: Visualization of a vertex i and its dual $\star i$.

rally compatible with additional energy terms for normal alignment, landmarks, smoothing, and nonpenetration.

1.3 The role of the ARAP term

The retargeting energy is driven by two competing geometric requirements. On one hand, the conforming surface should move into contact or near-contact with the target. On the other hand, it should preserve as much of its original structure as the task permits. A distance or clearance term addresses the first requirement, but by itself it says nothing about which features of the conforming mesh should be preserved during the deformation. In the garment example, an optimization driven only by contact may move sleeve openings, distort seams, and alter proportions even if it succeeds in eliminating interpenetration.

This is the role of the ARAP term. In this thesis, “rigidity” means near-isometry in the ARAP sense: local edge neighborhoods should deform approximately by rotations, so lengths and angles are preserved as much as the target constraints allow. This is not a full mechanical rigidity

model and should not be interpreted as a bending- or shear-resistance law. In our framework this near-isometry term is combined with the competing SDF-driven terms to produce favorable deformations. Additional normal, landmark, smoothing, and barrier terms may then be added to control particular aspects of the deformation. The result is a framework in which contact, preservation, alignment, and nonpenetration can be balanced against one another in a controlled way.

1.4 Main results and organization of the thesis

The thesis has three principal contributions.

First, it develops a customizable SDF-plus-energies framework for the class of problems described above. The target is represented by an ADSDF, the conforming mesh carries both Euclidean and manifold-valued data, and the optimization is carried out intrinsically on the product space $(\mathbb{R}^3)^N \times (S^3)^N$. Relative to the classical ARAP pipeline, the essential difference is that positions and rotations are updated within the same intrinsic descent procedure rather than by an interleaved local-global iteration. The resulting framework is compatible with additional energy terms for normal alignment, landmarks, smoothing, and nonpenetration.

Second, it establishes a local theorem relating the two most important higher-order ADSDF constructions: Taylor samples and least-squares (LSQ) samples. The ADSDF paper of Valasek and Bán suggests that LSQ samples converge to Taylor samples as the fine-grid extent shrinks [2]. One contribution of this thesis is to state and prove the local comparison theorem needed to make that statement precise.

Theorem 1.4.1 (Informal LSQ-to-Taylor principle).

On regular regions, degree- k LSQ ADSDF samples approximate the target signed distance function to the same order as degree- k Taylor ADSDF samples. Therefore the filtered LSQ ADSDF inherits the convergence order of the corresponding filtered Taylor ADSDF.

The formal statements and proofs are given in Chapter 3 in the form of Theorem 3.3.2 and Corollary 3.3.2.

For signed distance fields, the hypothesis on K is local by necessity. The field is not globally smooth: it fails to be C^1 on the cut locus, where closest points are not unique. Thus the theorem applies only on regions where the target field is smooth. Even so, it justifies the connection between LSQ ADSDFs and Taylor ADSDFs in practice.

Third, the thesis develops a subdivision result for manifold-valued data. Classical subdivision schemes such as Catmull–Clark and Loop refine Euclidean control meshes by repeated weighted averaging [9–11]. In our framework the per-vertex quaternion data must also be refined, and the correct replacement for affine averaging is a Karcher mean. This leads to a manifold-valued analogue of these schemes and, under standard proximity hypotheses, to a convergence result of the following form [12–14].

Theorem 1.4.2 (Informal subdivision theorem).

Suppose a subdivision rule on a triangular or quadrilateral mesh is obtained by replacing the affine averages in a classical linear scheme with the corresponding Karcher means on a Riemannian manifold. Under suitable small-data and proximity hypotheses, the refined control meshes converge to a continuous manifold-valued limit surface.

The formal statement and proof are given in Chapter 5 in the form of Theorem 5.6.1.

This manifold-valued subdivision material is a non-essential companion result to the central retargeting algorithm, but it does provide a starting point for developing more advanced multi-resolution strategies in future work and it is mathematically interesting in its own right. A basic way to apply our subdivision result is a coarse-to-fine optimization strategy: solve a global fitting problem at low resolution, then subdivide and continue only when added detail is needed.

At a high level, the remainder of the thesis is organized as follows.

- Chapter 2 gives the signed-distance preliminaries needed for retargeting and then develops the energy terms themselves.
- Chapter 3 develops algebraic discrete signed distance fields, explains their construction and filtering, and proves the local LSQ-to-Taylor comparison theorem.

- Chapter 4 develops the intrinsic optimization framework, including Riemannian gradients, Lie-group structure, quaternion formulas, and the Gauss–Newton viewpoint used later.
- Chapter 5 presents the companion subdivision result for manifold-valued data and explains its role in coarse-to-fine refinement.
- Chapter 6 records the illustrative applications and experimental formulations used to demonstrate the framework in practice.
- Chapter 7 concludes with a summary of contributions, limitations, and future directions.

The unifying theme is that implicit target geometry, manifold-valued deformation data, and customizable energies fit together naturally. The ADSDF supplies local geometric information about the target. The manifold-valued data encode rotational structure in a way compatible with intrinsic optimization. The energy framework makes it possible to ask for deformations that preserve design intent while simultaneously promoting contact, normal alignment, and other geometric properties.

The library developed to carry out these computations is available at

<https://github.com/tristanneighbors/Energy-Framework-Contact-Aware-Mesh-Deformation>.

Chapter 2

Retargeting Energies and SDFs

This chapter defines the energy terms that drive the retargeting problems considered in this thesis. The guiding example is a garment or accessory that has already been designed for one character and must be adapted to another. A shirt authored for a base humanoid, for example, cannot simply be placed onto a more muscular body and left unchanged: the geometry may intersect the body mesh unrealistically—a phenomenon commonly called **clipping** in computer graphics—and even when intersection is avoided, cuffs, collars, seams, and other designed features may drift into implausible positions. The problem is therefore not merely to bring one surface into contact with another, but to deform the source mesh in a way that respects the geometric relationships that make the original asset look intentional.

Our approach is to represent the target surface by a signed distance function and to formulate retargeting as the minimization of a customizable energy. This chapter focuses on the idea of designing these energies. We first recall the small amount of smooth SDF geometry that we’ll need. We then describe the basic SDF-driven terms, the near-isometric term inspired by **as-rigid-as-possible** (ARAP) modeling, and a short list of auxiliary terms that can be added without changing the overall structure of the method. The practical questions—how to query the target surface discretely with accurate derivatives, and how to optimize a function of both Euclidean and manifold-valued variables—are deferred to the next two chapters.

2.1 Signed-distance preliminaries

We use the symbol $\mathcal{T}$ for the fixed **target surface**. The surface to be deformed toward $\mathcal{T}$ will be denoted by $\mathcal{C}$ and will be called the **conforming surface**. Throughout the thesis, $\mathcal{C}$ is represented as a mesh. The target surface will first be described ideally by a smooth signed distance function Φ_{SDF} and later, in Chapter 3, by a discrete approximation that can be evaluated efficiently in practice.

2.1.1 Signed distance functions, regular points, and the cut locus

Definition 2.1.1.

Let $\Omega \subset \mathbb{R}^3$ have nonempty boundary $\partial\Omega$. The signed distance function of Ω is the map $\Phi_{\text{SDF}} : \mathbb{R}^3 \rightarrow \mathbb{R}$ defined by

$$\Phi_{\text{SDF}}(\mathbf{x}) = \begin{cases} -\text{dist}(\mathbf{x}, \partial\Omega), & \mathbf{x} \in \Omega, \\ \text{dist}(\mathbf{x}, \partial\Omega), & \mathbf{x} \notin \Omega. \end{cases}$$

Its zero level set is the boundary:

$$\partial\Omega = \{\mathbf{x} \in \mathbb{R}^3 : \Phi_{\text{SDF}}(\mathbf{x}) = 0\}.$$

For the energy terms below, the important point is that a signed distance function packages several kinds of geometric information into a single scalar field. The value $\Phi_{\text{SDF}}(\mathbf{x})$ records distance to the target surface together with an inside/outside sign. At points where Φ_{SDF} is differentiable, its gradient gives a unit normal direction. At points where the Hessian exists, it records curvature information for the level set passing through $\mathbf{x}$.

Signed distance functions are continuous and 1-Lipschitz, but they are not globally smooth. In a neighborhood of a point with a unique closest point on a smooth target surface, the signed distance function behaves as expected and is as smooth as the target. We call such points **regular points**, and regions containing only regular points **regular regions**. At points where the closest point is not unique, differentiability can fail; this set is referred to as the **cut locus**.

We recall the standard SDF facts needed for our discussion of retargeting energies. Proofs are omitted but they can be found in standard references [3].

Proposition 2.1.1.

Let $\mathbf{x}$ be a regular point of Φ_{SDF} , and let

$$\mathcal{S}_c = \{y \in \mathbb{R}^3 : \Phi_{\text{SDF}}(y) = c\}, \quad c = \Phi_{\text{SDF}}(\mathbf{x}),$$

be the level set through $\mathbf{x}$. Then

$$\underbrace{\|\nabla \Phi_{\text{SDF}}(\mathbf{x})\|}_{\text{unit length}} = 1, \quad \underbrace{\nabla \Phi_{\text{SDF}}(\mathbf{x})}_{\text{normal direction}} \perp T_{\mathbf{x}}\mathcal{S}_c.$$

2.1.2 The Hessian and curvature of level sets

When $\nabla^2 \Phi_{\text{SDF}}(\mathbf{x})$ exists, it describes the second-order geometry of the level set through $\mathbf{x}$. On the target surface itself this means the zero level set; away from contact it means an offset surface.

Proposition 2.1.2.

Let $\mathbf{x}$ be a regular point of Φ_{SDF} , and let $\mathcal{S}_c$ be the level set around $\mathbf{x}$ as above. Then $\nabla \Phi_{\text{SDF}}(\mathbf{x})$ is an eigenvector of $\nabla^2 \Phi_{\text{SDF}}(\mathbf{x})$ with eigenvalue 0. The other two eigenvectors lie in $T_{\mathbf{x}}\mathcal{S}_c$, and their eigenvalues are the principal curvatures of $\mathcal{S}_c$ at $\mathbf{x}$. Consequently, if these principal curvatures are denoted by κ_1, κ_2 , then

$$\underbrace{\text{tr}(\nabla^2 \Phi_{\text{SDF}}(\mathbf{x}))}_{\text{mean curvature of } \mathcal{S}_c} = \kappa_1 + \kappa_2$$

and

$$\underbrace{\frac{1}{2} \left(\text{tr}(\nabla^2 \Phi_{\text{SDF}}(\mathbf{x}))^2 - \text{tr}(\nabla^2 \Phi_{\text{SDF}}(\mathbf{x}))^2 \right)}_{\text{Gaussian curvature of } \mathcal{S}_c} = \kappa_1 \kappa_2.$$

These identities are standard consequences of the spectral theorem together with the Cayley–Hamilton relation for the elementary symmetric polynomials of the eigenvalues [3]. More explic-

itly, at a regular point the Hessian has eigenvalues $0, \kappa_1, \kappa_2$. Hence

$$\text{tr}(\nabla^2 \Phi_{\text{SDF}}) = \kappa_1 + \kappa_2, \quad \text{tr}(\nabla^2 \Phi_{\text{SDF}}^2) = \kappa_1^2 + \kappa_2^2,$$

and therefore

$$\frac{1}{2} (\text{tr}(\nabla^2 \Phi_{\text{SDF}})^2 - \text{tr}(\nabla^2 \Phi_{\text{SDF}}^2)) = \frac{1}{2} ((\kappa_1 + \kappa_2)^2 - (\kappa_1^2 + \kappa_2^2)) = \kappa_1 \kappa_2.$$

The chapter does not need more smooth SDF theory than this. Appendix A collects additional reference material: a compact characterization of signed distance functions, several exact formulas for common primitives, and explicit gradients and Hessians that are useful for testing implementations.

2.2 The retargeting energy at a glance

We now turn from smooth target-side geometry to the discrete variables that will actually be optimized. Let $\mathcal{V} = \{1, \dots, N\}$ denote the vertices of the conforming mesh. The vertex positions are collected into

$$\mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_N) \in (\mathbb{R}^3)^N.$$

For the near-isometric term we will also attach to each vertex an auxiliary rotation variable

$$\mathbf{Q} = (q_1, \dots, q_N) \in (S^3)^N,$$

interpreted as unit quaternions. The one-ring of a vertex i is denoted by $\mathcal{N}_1(i)$, and $\text{area}(\star i)$ denotes the dual area associated to i .

The energy is customizable. Some applications may be concerned with normal alignments and others with conformity to certain landmarks. Some may be concerned with multiple geometric features simultaneously. One designs the energy to suit the specific application. The one-ring neighborhood $\mathcal{N}_1(i)$ is the set of vertices adjacent to vertex i by an edge.




Figure 2.1: One-ring neighborhood $\mathcal{N}_1(i)$ (yellow) of vertex i (teal).

For garment retargeting the basic ingredients are a term that responds to the target SDF and a term that preserves the designed mesh as nearly isometric as possible. Additional terms can then enforce further geometric preferences such as normal alignment, semantic landmarks, or auxiliary smoothness.

A convenient schematic form is

$$\begin{aligned}
 \mathcal{E}_{\text{total}}(\mathbf{P}, \mathbf{Q}) = & \underbrace{T_{\text{dist}}\mathcal{E}_{\text{dist}}(\mathbf{P}) + T_{\text{clear}}\mathcal{E}_{\text{clear}}(\mathbf{P}) + T_{\text{norm}}\mathcal{E}_{\text{norm}}(\mathbf{P}, \mathbf{Q})}_{\text{driven directly by SDF queries}} \\
 & + \underbrace{T_{\text{ARAP}}\mathcal{E}_{\text{ARAP}}(\mathbf{P}, \mathbf{Q})}_{\text{promotes near-isometry}} \\
 & + \underbrace{T_{\text{land}}\mathcal{E}_{\text{land}}(\mathbf{P}) + T_{\text{smooth}}\mathcal{E}_{\text{smooth}}(\mathbf{P}, \mathbf{Q}) + \dots}_{\text{auxiliary, application-specific terms}}.
 \end{aligned}$$

The coefficients $T_{\text{dist}}, T_{\text{clear}}, \dots$ are energy weights. They may be chosen by hand, prescribed regionwise, or scheduled automatically later in the optimization. Their role is only to modulate the relative influence of the component energies; they do not change the basic structure of the method.

2.3 SDF-driven terms

This section records the target-side terms that follow immediately from the preliminaries above. They are the terms that depend explicitly on values, normals, or possibly Hessians of the target SDF.

2.3.1 Distance, offset, and clearance

The simplest fitting term asks the conforming surface to approach a prescribed level set of the target SDF. If the desired offset is the scalar field d_i defined at the vertices of the conforming mesh, one natural choice is

$$\mathcal{E}_{\text{dist}}(\mathbf{P}) = \sum_{i \in \mathcal{V}} \underbrace{(\Phi_{\text{SDF}}(\mathbf{p}_i) - d_i)^2}_{\text{distance to desired level set}} \text{area}(\star i).$$

The special case $d_i \equiv 0$ asks the mesh to approach the target surface itself. More generally, a positive offset can be used to impose a desired stand-off distance. This is often more appropriate for clothing and accessories than exact contact, since the goal is not to glue the source mesh to the body but to make it fit plausibly around it.

A two-sided distance term is appropriate when the conforming surface is intended to lie near a prescribed offset level set. It is less appropriate for loose garments or for regions where contact should occur only on one side, because it can pull the mesh too tightly around the target, especially near concavities. In such cases the distance term should be replaced or supplemented by a one-sided clearance term, landmarks, and ARAP regularization so that the mesh avoids penetration without being forced to cling to the target everywhere.

For many retargeting tasks, contact alone is not enough. One often wants a one-sided penalty that forbids or strongly discourages interpenetration. A simple choice is

$$\mathcal{E}_{\text{clear}}(\mathbf{P}) = \sum_{i \in \mathcal{V}} \underbrace{\max(0, d_i - \Phi_{\text{SDF}}(\mathbf{p}_i))^2}_{\text{penalty for violating the clearance}} \text{area}(\star i),$$

where $d_i \geq 0$ is again the desired offset. This term vanishes whenever the vertex lies on or outside the desired offset level set and becomes active only when the clearance is violated. In practice

one may prefer smoother barrier-type penalties. The formula above is a simple example of a non-smooth penalty that is easy to implement and works well enough in some applications.

2.3.2 Normal alignment

Once the conforming surface is near contact, normal information becomes meaningful. Let

$$\mathbf{N} = (\mathbf{n}_1, \dots, \mathbf{n}_N), \quad \mathbf{n}_i \in \mathbb{R}^3,$$

denote a reference normal field attached to the conforming mesh. A simple alignment energy is

$$\mathcal{E}_{\text{norm}}(\mathbf{P}, \mathbf{Q}) = \sum_{i \in \mathcal{V}} \chi_i \underbrace{\left(1 - (q_i \mathbf{n}_i q_i^*) \cdot \nabla \Phi_{\text{SDF}}(\mathbf{p}_i)\right)}_{\text{misalignment with target normal}} \text{area}(\star i),$$

where the weights $\chi_i \in [0, 1]$ activate the term only where it is intended to matter.

Two comments are important here. First, the unit quaternion q_i is meant to encode the local rotational part of the deformation near vertex i . After identifying vectors in $\mathbb{R}^3$ with pure imaginary quaternions, the map

$$\mathbf{n}_i \longmapsto q_i \mathbf{n}_i q_i^*$$

is the rotation of $\mathbf{n}_i$ determined by q_i . Thus $q_i \mathbf{n}_i q_i^*$ is the discrete analogue of transporting a reference normal by the local rotational component of the deformation. The quaternion conventions used here, including the action $v \mapsto qvq^*$, are reviewed in Appendix B. Second, this normal-alignment term should usually be applied only near intended contact regions. Away from contact, the gradient of the target SDF is the normal of an offset level set through the query point, and it may be geometrically unreasonable to force that direction to match a garment or accessory normal everywhere.

2.4 Near-isometry and the ARAP term

The central design-preserving term is the one inspired by ARAP modeling [1]. The original ARAP method was introduced for interactive surface deformation, not for garment retargeting, but the geometric idea carries over directly: when exact rigidity is impossible, one should still prefer deformations that are as close to locally rigid motions as the target constraints allow.

2.4.1 Smooth rigidity and its rotational part

Let $X_0 : \mathcal{C} \rightarrow \mathbb{R}^3$ denote the initial embedding of the conforming surface and let $X : \mathcal{C} \rightarrow \mathbb{R}^3$ denote a deformed embedding. The induced deformation map is

$$f = X \circ X_0^{-1}.$$

At a point $p \in \mathcal{C}$, the differential df_p records the tangential part of the deformation. If we adjoin the unit normal of the deformed surface, we obtain a 3×3 matrix

$$D_p = \left[df_p \mid n(p) \right].$$

After choosing compatible oriented tangent bases, this matrix admits a polar decomposition

$$D_p = R_p S_p, \quad R_p \in \text{SO}(3), \quad S_p \text{ symmetric positive definite.}$$

The factor R_p is the closest rotational part of the local deformation, while S_p is a symmetric positive-definite stretch tensor. The tensor S_p records the non-rotational part of the local deformation, including dilation and anisotropic or shear-like stretch. A pure local rotation is characterized by $S_p = I$. Equivalently, the deformation preserves lengths and angles to first order. This is the sense in which rigidity appears here.

The same decomposition explains why normal alignment naturally couples to the rotational variable. If $n_0(p)$ denotes the initial normal, then

$$\underbrace{R_p n_0(p)}_{\text{rotated reference normal}} = \underbrace{n(p)}_{\text{deformed normal}}.$$

The rotational factor is therefore the object that simultaneously describes how tangent directions and normals move under the deformation.

2.4.2 The discrete near-isometric energy

In the discrete setting we do not manipulate the full tensor df_p directly. Instead, we compare each deformed mesh edge with the corresponding edge in the reference mesh after rotating the reference edge by a per-vertex auxiliary variable. Let

$$\mathbf{P}^{(0)} = (\mathbf{p}_1^{(0)}, \dots, \mathbf{p}_N^{(0)})$$

denote the reference positions of the conforming mesh. The basic discrete near-isometric term is

$$\mathcal{E}_{\text{ARAP}}(\mathbf{P}, \mathbf{Q}) = \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{N}_1(i)} \mathbf{w}_{ij} \underbrace{\left\| q_i(\mathbf{p}_i^{(0)} - \mathbf{p}_j^{(0)}) q_i^* - (\mathbf{p}_i - \mathbf{p}_j) \right\|}_{\text{deformed edge minus rotated reference edge}}^2 \text{area}(\star i).$$

Here $\mathbf{w}_{ij} \geq 0$ are edge weights, identical to those appearing in the traditional ARAP method. The role of the term is easy to read from the formula: if the deformed edge $(\mathbf{p}_i - \mathbf{p}_j)$ agrees with the rotated reference edge at every incident pair (i, j) , then the deformation is locally consistent with a rigid motion. The farther these quantities drift apart, the larger the penalty.

This term promotes rigidity but does not strictly constrain the deformation to being rigid. Exact rigidity is generally incompatible with the retargeting problem. A shirt that must fit a larger torso or a wider shoulder span cannot remain exactly isometric. The role of $\mathcal{E}_{\text{ARAP}}$ is therefore to preserve the original design as much as the target-side terms permit, not to prevent all geometric change.

2.4.3 Relation to classical ARAP

The energy above is close in spirit to the classical ARAP energy, but the optimization viewpoint is different. In the original local/global formulation, rotational variables and positional variables are updated in alternating subproblems [1]. Here the rotational variables remain part of the same optimization state as the vertex positions. The next chapters will describe how this coupled objective is handled: Chapter 3 supplies a target representation with accurate derivatives, and a later intrinsic-optimization chapter (Chapter 4) develops the framework for the mixed Euclidean/manifold-valued variables.

The distinction matters because the customizable energy viewpoint is one of the main advantages of the framework presented in this thesis. Once the objective is written on the joint state (P, Q) , one can add or remove other terms—clearance, landmarks, smoothing, and so forth—without redesigning the overall modeling framework.

2.5 Auxiliary terms

Once the target-side terms and the near-isometric term are in place, additional energies can encode task-specific preferences without changing the overall form of the problem.

2.5.1 Landmarks

A **landmark** is a designated feature of the source mesh that should remain near a corresponding feature on the target. In a garment example, landmarks might represent cuffs, collar points, seam endpoints, or a hem curve. The simplest point-based version is

$$\mathcal{E}_{\text{land}}(P) = \sum_{\ell=1}^m \omega_{\ell} \|\mathbf{p}_{i_{\ell}} - \mathbf{x}_{\ell}^{\star}\|^2,$$

where i_{ℓ} indexes a vertex or representative point on the conforming mesh, $\mathbf{x}_{\ell}^{\star}$ is the prescribed target feature, and ω_{ℓ} controls how strongly that correspondence is enforced. More generally, one could use curve-based or region-based landmarks, and the energy could be designed to enforce tangential as well as positional alignment.

2.5.2 Auxiliary smoothing or regularization

A second family of auxiliary terms penalizes rough or oscillatory deformations. One option is to smooth the displacement field directly. Another, more natural in the present framework, is to smooth the auxiliary rotational data. If the rotations are represented by unit quaternions, one convenient choice is

$$\mathcal{E}_{\text{smooth}}(\mathbf{Q}) = \sum_{i \in \mathcal{V}} \frac{1}{|\mathcal{N}_1(i)|} \sum_{j \in \mathcal{N}_1(i)} \frac{1}{2} \|\log(q_i^* q_j)\|^2 \text{area}(\star i).$$

The logarithm here is the Riemannian logarithm on S^3 , so the term measures squared geodesic discrepancy between neighboring rotations. Chapter 4 will review the small amount of manifold geometry needed to interpret this expression properly. For this chapter, the important point is simply that such a term is available if the deformation requires additional regularity [15].

2.5.3 Energy weights and regionwise compliance

The coefficients

$$T_{\text{dist}}, T_{\text{clear}}, T_{\text{norm}}, T_{\text{ARAP}}, T_{\text{land}}, T_{\text{smooth}}, \dots$$

should be viewed as compliance controls. They may be constant, chosen interactively by a designer, or allowed to vary by region. For example, a garment panel near the torso may be permitted to deviate more from isometry than a cuff or collar. Likewise, landmarks may be enforced strongly only in a few semantic regions. The later optimization chapter will separate clearly between the fixed-weight inner problem and any outer continuation strategy. They are called weights throughout the discussion below; the notation T_* is only a compact way of indexing the residual blocks.

2.6 What remains before optimization is possible

The chapter has defined the ideal target-side quantities and the energy terms that use them, but two substantial ingredients are still missing.

First, the target surface must be represented discretely in a way that preserves the value of the signed distance formulation. The energies above require repeated access not only to a scalar field value but also to its derivatives. Finite-difference derivatives of a sampled scalar grid are often too noisy and stencil-dependent for this role. The next chapter therefore turns to algebraic discrete signed distance fields, whose local polynomial structure provides analytic derivatives together with a useful local approximation theory [2].

Second, the state variable $(\mathbf{P}, \mathbf{Q})$ mixes Euclidean and manifold-valued data. The positional variables live in $(\mathbb{R}^3)^N$, but the rotational variables live on $(S^3)^N$ (and more generally could live on another Riemannian manifold). To minimize the energy we therefore need a framework that differentiates and updates both parts coherently. Chapter 4 develops that framework in terms of intrinsic optimization on product spaces.

Chapter 3

Algebraic Discrete Signed Distance Fields

Chapter 2 described the energy terms in terms of an ideal signed distance function Φ_{SDF} . To use those energies in practice, one needs a discrete representation of the target surface that supports repeated value, gradient, and sometimes Hessian queries. This chapter develops the representation used for that purpose: algebraic discrete signed distance fields (ADSDFs) [2].

The central point is that ordinary sampled signed distance fields store only scalar values and leave derivative information to finite differences. For optimization, that is often too weak. ADSDFs instead store local polynomial data, so that derivatives are available analytically from the representation itself. The chapter first reviews the passage from ordinary DSDFs to higher-order algebraic sample fields, then explains the practical issues of filtering and basis choice, and finally culminates in a local theorem that relates least-squares ADSDFs to their Taylor counterparts on regular regions.

Subsections 3.1 and 3.2 are mostly a reformulation of the material developed by Valasek and Bán in their ADSDF paper [2] with the addition of a more rigorous discussion of numerical stability. Subsection 3.3 covers the formal version of the informally stated LSQ-to-Taylor comparison theorem from Chapter 1.

3.1 From DSDFs to ADSDFs

3.1.1 Discrete signed distance functions as a baseline

A continuous signed distance function is rarely available in closed form for a production mesh, so one typically works with a discrete approximation. The most common construction is to sample distance values at the vertices of a regular grid and extend those samples by trilinear interpolation. More generally, one may work on any cellular decomposition of a domain $\mathcal{D} \subset \mathbb{R}^3$.

Let $\mathcal{I} = \{1, \dots, M\}$ index the sample sites, let $\mathbf{x}_i \in \mathbb{R}^3$ denote the sample positions, and let

$$d_i = \Phi_{\text{SDF}}(\mathbf{x}_i)$$

be the stored sample values. If $\lambda_i(\mathbf{x})$ are barycentric basis functions supported on the cells incident to sample i , then the induced discrete signed distance field is

$$\hat{\Phi}_{\text{SDF}}(\mathbf{x}) = \sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) d_i.$$

When the target begins as a triangle mesh, the sample value d_i is computed by a closest-feature query against the mesh. The unsigned distance is the minimum Euclidean distance to the target triangles, edges, and vertices. The sign is assigned by a ray-parity test against the closed target mesh: points with an odd number of ray intersections are treated as inside and points with an even number are treated as outside. These signed samples are the data from which the local least-squares polynomial samples are fitted.

This baseline DSDF is simple, memory efficient, and often sufficient when only values are needed [2, 16]. Its weakness is equally clear: derivatives come only from differentiating the interpolant, or from applying finite differences to the reconstructed field, neither of which faithfully stores higher-order information from the original geometry.

ADSDFs address this by replacing scalar samples with local polynomial samples.

3.1.2 Higher-order sample fields and Taylor ADSDFs

The usual DSDF corresponds to storing scalar values. An ADSDF stores instead a local polynomial, so that each sample describes the behavior of the target function on a neighborhood rather than at a single point.

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a sufficiently smooth function and let $\lambda_i(\mathbf{x})$ be nonnegative barycentric weights with $\sum_i \lambda_i(\mathbf{x}) = 1$. Suppose each sample point $\mathbf{x}_i$ carries a local approximation $f_i(\mathbf{x})$ to $f(\mathbf{x})$.

Theorem 3.1.1 (Convex-combination error bound).

Assume that the local approximants satisfy

$$|f(\mathbf{x}) - f_i(\mathbf{x})| \leq R_i(\mathbf{x})$$

on the cells where λ_i is supported. Then the reconstructed field

$$\hat{f}(\mathbf{x}) = \sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) f_i(\mathbf{x})$$

satisfies

$$|f(\mathbf{x}) - \hat{f}(\mathbf{x})| \leq \sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) R_i(\mathbf{x}) \leq \max_{i \in \mathcal{I}} \max_{\mathbf{x} \in \text{cells}(i)} R_i(\mathbf{x}).$$

Proof. Since the λ_i form a nonnegative partition of unity,

$$|f(\mathbf{x}) - \hat{f}(\mathbf{x})| = \left| \sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) (f(\mathbf{x}) - f_i(\mathbf{x})) \right| \leq \sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) |f(\mathbf{x}) - f_i(\mathbf{x})|.$$

Substituting the bounds R_i gives the claim. □

This theorem is the basic local mechanism behind Taylor ADSDFs. If the local polynomial samples are accurate on each cell, then any convex barycentric blend of them inherits the same worst-case order.

We will use standard multi-index notation. For $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{N}^n$, write

$$|\alpha| = \alpha_1 + \dots + \alpha_n, \quad \alpha! = \alpha_1! \dots \alpha_n!, \quad (\mathbf{x} - \mathbf{x}_i)^\alpha = \prod_{m=1}^n (x_m - x_{i,m})^{\alpha_m},$$

and

$$\partial^\alpha f = \frac{\partial^{|\alpha|} f}{\partial x_1^{\alpha_1} \dots \partial x_n^{\alpha_n}}.$$

The degree- k Taylor polynomial of f about $\mathbf{x}_i$ is

$$T_{f;\mathbf{x}_i}^k(\mathbf{x}) = \sum_{|\alpha| \leq k} \frac{\partial^\alpha f(\mathbf{x}_i)}{\alpha!} (\mathbf{x} - \mathbf{x}_i)^\alpha.$$

Corollary 3.1.1 (Taylor ADSDF error bound).

Let the local sample at $\mathbf{x}_i$ be the degree- k Taylor polynomial $T_{f;\mathbf{x}_i}^k$. Assume $f \in C^{k+1}$ on a neighborhood of each cell and define

$$h_i = \max_{\mathbf{x} \in \text{cells}(i)} \|\mathbf{x} - \mathbf{x}_i\|_1, \quad M_i = \max_{|\alpha|=k+1} \sup_{\|\vec{\mathbf{h}}\|_1 \leq h_i} |\partial^\alpha f(\mathbf{x}_i + \vec{\mathbf{h}})|.$$

Then the Taylor ADSDF

$$\hat{f}_k(\mathbf{x}) = \sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) T_{f;\mathbf{x}_i}^k(\mathbf{x})$$

satisfies

$$|f(\mathbf{x}) - \hat{f}_k(\mathbf{x})| \leq \max_{i \in \mathcal{I}} \frac{M_i}{(k+1)!} h_i^{k+1}.$$

Proof. Taylor's theorem with Lagrange remainder gives

$$f(\mathbf{x}) = T_{f;\mathbf{x}_i}^k(\mathbf{x}) + \underbrace{\sum_{|\alpha|=k+1} \frac{\partial^\alpha f(\mathbf{x}_i + \theta(\mathbf{x} - \mathbf{x}_i))}{\alpha!} (\mathbf{x} - \mathbf{x}_i)^\alpha}_{R_{f;\mathbf{x}_i}^k(\mathbf{x})}$$

for some $\theta \in (0, 1)$. Hence

$$|R_{f;\mathbf{x}_i}^k(\mathbf{x})| \leq M_i \sum_{|\alpha|=k+1} \frac{|\mathbf{x} - \mathbf{x}_i|^\alpha}{\alpha!} = \frac{M_i}{(k+1)!} \|\mathbf{x} - \mathbf{x}_i\|_1^{k+1} \leq \frac{M_i}{(k+1)!} h_i^{k+1},$$

where the multinomial theorem is used in the equality. The result now follows from Theorem 3.1.1. □

A useful special case is the order-0 field. Then each sample is merely the constant value $f(\mathbf{x}_i)$, and the error estimate reduces to the familiar statement that zeroth-order distance samples describe the function correctly only to first order away from the sample site.

3.1.3 Why the interpolants matter: derivatives and Hermite-type endpoint conditions

The main reason to prefer ADSDFs is not only that they improve value accuracy, but that they permit analytic derivatives of the reconstructed field. However, the derivatives of the field are not determined solely by the sample polynomials. They also depend on the blending functions.

Indeed, for

$$\hat{f}(\mathbf{x}) = \sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) f_i(\mathbf{x}),$$

we do *not* have

$$\partial^\alpha \hat{f} = \sum_i \lambda_i \partial^\alpha f_i$$

automatically. Product-rule terms appear:

$$\nabla \hat{f}(\mathbf{x}) = \underbrace{\sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) \nabla f_i(\mathbf{x})}_{\text{desired term}} + \underbrace{\sum_{i \in \mathcal{I}} (\nabla \lambda_i(\mathbf{x})) f_i(\mathbf{x})}_{\text{blending term}},$$

and higher derivatives produce even more such terms. The problem is therefore not simply to interpolate values, but to choose blending functions whose derivatives vanish in the right places.

It is helpful to recall what Hermite interpolation is for. Ordinary interpolation matches function values at specified points. Hermite interpolation matches values *and* prescribed derivatives. In one dimension, this means that if two sample functions are blended across an interval, the blending function must hit the correct endpoint values and kill enough endpoint derivatives. That is exactly the role played here by the Hermite-type smoothing polynomials.

Consider a one-dimensional cell $[0, 1]$ with endpoint sample functions f_0, f_1 . Write

$$\hat{f}(t) = (1 - h(t))f_0(t) + h(t)f_1(t).$$

To reproduce derivatives up to order k at the endpoints, it is enough to require

$$h(0) = 0, \quad h(1) = 1, \quad h^{(m)}(0) = h^{(m)}(1) = 0 \quad \text{for } m = 1, \dots, k.$$

The first few solutions are

$$h_1(t) = t, \quad h_3(t) = -2t^3 + 3t^2, \quad h_5(t) = 6t^5 - 15t^4 + 10t^3,$$

$$h_7(t) = -20t^7 + 70t^6 - 84t^5 + 35t^4.$$

These are exactly the linear, cubic, quintic, and septic Hermite-type blending functions used in [2].

For a cubical cell with normalized coordinates $(u, v, w) \in [0, 1]^3$, define

$$\vartheta_0(t) = 1 - h(t), \quad \vartheta_1(t) = h(t),$$

and for each vertex $(a, b, c) \in \{0, 1\}^3$ set

$$\lambda_{abc}(u, v, w) = \vartheta_a(u)\vartheta_b(v)\vartheta_c(w).$$

Then the reconstructed field on the cell takes the tensor-product form

$$\hat{f}(u, v, w) = \sum_{a,b,c \in \{0,1\}} \lambda_{abc}(u, v, w) f_{abc}(u, v, w).$$

Proposition 3.1.1 (Derivative reproduction at cell vertices).

Let h satisfy the endpoint conditions above up to order k , and let λ_{abc} be the associated tensor-product weights. Then for every cell vertex $\mathbf{x}_{abc}$ and every multi-index α with $|\alpha| \leq k$,

$$\partial^\alpha \hat{f}(\mathbf{x}_{abc}) = \partial^\alpha f_{abc}(\mathbf{x}_{abc}).$$

Proof. Every term in $\partial^\alpha \hat{f}$ in which a derivative falls on an “incorrect” tensor factor acquires a derivative of ϑ_0 or ϑ_1 evaluated at an endpoint. By construction those derivatives vanish up to order k . The only surviving contribution is the one in which the weight equals 1 at the active vertex and all differentiated weight factors vanish. $\square$

After Proposition 3.1.1, the derivative version of Corollary 3.1.1 is essentially the same argument as before.

Corollary 3.1.2 (Derivative error for Taylor ADSDFs).

Let $\hat{f}_k$ be the Taylor ADSDF from Corollary 3.1.1, and assume the blending functions reproduce derivatives through order k at cell vertices in the sense of Proposition 3.1.1. Then for every multi-index β with $|\beta| \leq k$,

$$\partial^\beta \hat{f}_k$$

approximates $\partial^\beta f$ on each cell with order $k + 1 - |\beta|$.

Once the product-rule terms have been neutralized at the sample sites, the derivative statement is nothing more than Taylor’s theorem applied to $\partial^\beta f$ rather than to f itself.




Figure 3.1: One-dimensional Hermite-type blending functions.

3.2 Representing, filtering, and constructing ADSDFs in practice

By *filtering*, we mean the reconstruction step that blends local polynomial samples with the cell weights λ_i . In value form this is $\hat{f}(\mathbf{x}) = \sum_i \lambda_i(\mathbf{x}) f_i(\mathbf{x})$. In coefficient-pooling form, the same operation is performed by blending polynomial coefficients before evaluating the resulting polynomial.

3.2.1 Local coordinates, global monomials, and why normalization matters

A Taylor sample is most naturally written in centered coordinates:

$$f_i(\mathbf{x}) = \sum_{|\alpha| \leq k} c_{i,\alpha} (\mathbf{x} - \mathbf{x}_i)^\alpha.$$

This is the mathematically natural basis because Taylor's theorem is centered at $\mathbf{x}_i$. It is *not*, however, the computationally most convenient basis for filtering.

Suppose each local sample has already been translated into the global monomial basis

$$f_i(\mathbf{x}) = \sum_{|\alpha| \leq k} a_{i,\alpha} \mathbf{x}^\alpha.$$

Then the filtered ADSDF can be rewritten as

$$\underbrace{\sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) f_i(\mathbf{x})}_{\text{blend many samples}} = \sum_{|\alpha| \leq k} \underbrace{\left(\sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) a_{i,\alpha} \right)}_{\text{blend the coefficients instead}} \mathbf{x}^\alpha.$$

This is the key coefficient-pooling identity. Instead of evaluating several local polynomials and then blending the resulting values, one first interpolates the coefficients and then evaluates *one* polynomial. This is exactly the global-basis viewpoint that makes ADSDF queries GPU-friendly [2].

There is, however, an important caveat. The translation from centered coordinates $(\mathbf{x} - \mathbf{x}_i)^\alpha$ to global monomials $\mathbf{x}^\alpha$ is numerically stable only after the geometry has been normalized to a small canonical box, and in practice the natural choice is $[-1, 1]^3$. This point is implicit in [2]; here we make it explicit.

On $[-1, 1]^3$ every monomial $\mathbf{x}^\alpha$ with $|\alpha| \leq k$ satisfies $|\mathbf{x}^\alpha| \leq 1$. More generally, for each fixed multi-index β with $|\beta| \leq k$, the derivatives $\partial^\beta \mathbf{x}^\alpha$ are uniformly bounded on $[-1, 1]^3$ by constants depending only on k and β . Thus coefficient perturbations produce comparably sized *absolute* perturbations in the evaluated field, gradient, and Hessian.²

For signed distance fields this absolute control is geometrically meaningful. If $\widehat{\Phi} = \Phi + \delta\Phi$ and a regular point of the zero set is displaced by a normal amount δs , then the usual one-dimensional root-sensitivity formula along the normal direction gives

$$\delta s \approx -\frac{\delta\Phi}{\partial_n \Phi}.$$

²Relative value conditioning deteriorates near any root because the denominator $|p(\mathbf{x})|$ is small there. For signed distance fields, however, the geometric object of interest is the zero set itself, so absolute field error is the more relevant quantity near the interface.

At regular points of an exact SDF one has $\partial_n \Phi = \|\nabla \Phi\| = 1$, so the interface displacement is of the same order as the absolute field error. In this sense the geometry represented by an SDF is robust to small absolute perturbations, unlike a general polynomial whose derivative at a root may be arbitrarily small.

Proposition 3.2.1 (Quantitative stability of centered-to-global translation).

Let

$$\Pi_k(\mathbb{R}^n) = \left\{ p(y) = \sum_{|\alpha| \leq k} c_\alpha y^\alpha \right\}$$

denote the space of polynomials of total degree at most k . Fix k and suppose the sample positions lie in $\Omega \subset [-R, R]^n$. For each $\mathbf{x}_i \in \Omega$, let

$$\mathcal{T}_{\mathbf{x}_i} : \Pi_k(\mathbb{R}^n) \rightarrow \Pi_k(\mathbb{R}^n)$$

denote the translation map sending centered coordinates $(\mathbf{x} - \mathbf{x}_i)^\alpha$ to the global monomial basis $\mathbf{x}^\alpha$. Then, in coefficient ℓ^1 -norm,

$$\|\mathcal{T}_{\mathbf{x}_i}\|_1 \leq (1 + \|\mathbf{x}_i\|_\infty)^k, \quad \kappa_1(\mathcal{T}_{\mathbf{x}_i}) \leq (1 + \|\mathbf{x}_i\|_\infty)^{2k} \leq (1 + R)^{2k}.$$

In particular, on $[-1, 1]^n$ one has

$$\kappa_1(\mathcal{T}_{\mathbf{x}_i}) \leq 4^k.$$

For the low orders relevant here, this gives worst-case bounds 4, 16, 64 for $k = 1, 2, 3$. Consequently, centered coefficient estimates transfer to global-basis coefficient estimates with only a mild, explicitly controlled amplification on the canonical normalized box.

Proof. By the multivariate binomial formula,

$$(\mathbf{x} - \mathbf{x}_i)^\alpha = \sum_{\beta \leq \alpha} \binom{\alpha}{\beta} (-\mathbf{x}_i)^{\alpha - \beta} \mathbf{x}^\beta.$$

Thus the translation matrix has entries

$$(\mathcal{T}_{\mathbf{x}_i})_{\beta,\alpha} = \begin{cases} \binom{\alpha}{\beta} (-\mathbf{x}_i)^{\alpha-\beta}, & \beta \leq \alpha, \\ 0, & \text{otherwise.} \end{cases}$$

Using the coefficient ℓ^1 -norm, the column sum associated with α is

$$\sum_{\beta \leq \alpha} \left| \binom{\alpha}{\beta} (-\mathbf{x}_i)^{\alpha-\beta} \right| = \prod_{j=1}^n (1 + |x_{i,j}|)^{\alpha_j} \leq (1 + \|\mathbf{x}_i\|_\infty)^k.$$

Hence $\|\mathcal{T}_{\mathbf{x}_i}\|_1 \leq (1 + \|\mathbf{x}_i\|_\infty)^k$. Since translation back is translation by $-\mathbf{x}_i$, one has $\mathcal{T}_{\mathbf{x}_i}^{-1} = \mathcal{T}_{-\mathbf{x}_i}$ and therefore the same bound for the inverse. Multiplying the two estimates yields the stated condition-number bound. $\square$

The point of Proposition 3.2.1 is more specific than mere boundedness on an arbitrary bounded set. Every finite computation takes place on some bounded domain. What matters here is that the conditioning grows with the scale of that domain. Normalizing to $[-1, 1]^3$ keeps the global power basis well conditioned even for cubic samples, which is consistent with the empirical stability reported in [2] for normalized binary32 power-basis fields.

3.2.2 Cubical and simplicial filters

The cubical tensor-product case is the most convenient in code and is the case we will use by default. In normalized cell coordinates $(u, v, w) \in [0, 1]^3$, the tensor-product weights from the previous section are

$$\lambda_{abc}(u, v, w) = \vartheta_a(u) \vartheta_b(v) \vartheta_c(w), \quad (a, b, c) \in \{0, 1\}^3.$$

For $h = h_1$ this reduces to ordinary trilinear interpolation. For $h = h_3, h_5, h_7$ one obtains cubic, quintic, and septic derivative-preserving blending, respectively [2].

It is also useful to remember that the same idea works on simplices. Suppose the query point lies in a tetrahedron with barycentric coordinates (u_0, u_1, u_2, u_3) , so that $\sum_{m=0}^3 u_m = 1$. If h is a one-dimensional Hermite-type blend, one may define

$$\lambda_m(u) = h(u_m), \quad m = 1, 2, 3, \quad \lambda_0(u) = 1 - \lambda_1(u) - \lambda_2(u) - \lambda_3(u).$$

The same endpoint-derivative logic then shows that the reconstructed field reproduces the prescribed derivatives at the vertices. On a regular grid, one obtains a tetrahedral option by decomposing each cube into six tetrahedra. Valasek and Bán work out the resulting barycentric formulas explicitly; for the tetrahedron determined by the ordered coordinates $a \leq b \leq c$ one arrives at the compact weights

$$w_0 = 1 - c, \quad w_1 = a, \quad w_2 = c - b, \quad w_3 = b - a$$

after the appropriate choice of local tetrahedron [2]. We mention this option only to make the reader aware of its existence; the development in the thesis uses the cubical case.

3.2.3 Least-squares ADSDFs from fine stencils

Taylor ADSDFs assume that derivatives of the target function are available. This is realistic for analytic test geometries and for some procedural representations, but it is much less attractive when the target starts life as a mesh. Valasek and Bán therefore propose fitting local polynomials from additional distance samples rather than differentiating the geometry directly [2].

It is worth pausing to recall the general least-squares pattern. Given basis functions $\phi_1, \dots, \phi_m$, data values $y_1, \dots, y_N$, and positive weights $w_1, \dots, w_N$, weighted least squares seeks the coefficients that minimize the residual

$$\sum_{j=1}^N w_j \left| \sum_{\ell=1}^m a_\ell \phi_\ell(\xi_j) - y_j \right|^2$$




Figure 3.2: (Blue) A coarse sample point. (Red) A fine stencil of additional samples around it.

Let $V \in \mathbb{R}^{N \times m}$ be the design matrix $V_{j\ell} = \phi_\ell(\xi_j)$, and let $W = \text{diag}(w_1, \dots, w_N)$ be the diagonal matrix of sample weights. If the design matrix has full column rank, the solution is the familiar normal-equations formula

$$\mathbf{a} = (V^\top W V)^{-1} V^\top W \mathbf{y}.$$

In the ADSDF setting, one chooses a symmetric fine stencil around each coarse sample point and fits a degree- k polynomial to the sampled distance values on that stencil.

More concretely, let $\mathbf{x}_i$ be a coarse sample position and let

$$\Xi = \{\xi_1, \dots, \xi_N\} \subset \mathbb{R}^n$$

be a fixed normalized stencil about the origin. For a scale $\rho > 0$, the fine-grid samples are taken at

$$\mathbf{x}_i + \rho \xi_j, \quad j = 1, \dots, N.$$

Equivalently, if

$$g_{i,\rho}(y) = f(\mathbf{x}_i + \rho y),$$

then the least-squares sample $q_{i,\rho}$ is the degree- k polynomial minimizing

$$\sum_{j=1}^N w_j |q(\xi_j) - g_{i,\rho}(\xi_j)|^2.$$

Once that centered polynomial has been computed, one translates it to the global monomial basis and stores its coefficients. Since the normalized stencil is fixed, the pseudoinverse matrix can be precomputed once and reused at every coarse sample location [2].

From the perspective of later chapters, three practical observations matter most.

1. **Low order is the sweet spot.** Valasek and Bán report that orders 1 and 2 are the most attractive in practice. Order 3 is still usable in some settings, but the coefficient count is already much heavier, and higher orders rapidly become unattractive because of coefficient growth [2].
2. **The fine stencil acts like a local scale parameter.** Shrinking the stencil pushes the least-squares sample closer to the Taylor jet on regular regions, while enlarging it acts more like a local low-pass filter [2].
3. **Polynomial evaluation is not the conceptual bottleneck.** For the low degrees relevant here, Horner-type schemes are entirely adequate, and in practice hand-unrolled formulas for the polynomial and its derivatives are often simpler than a generic evaluation scheme.

degree k	coefficient count $\binom{3+k}{k}$
0	1
1	4
2	10
3	20
4	35

Table 3.1: Coefficient counts for total-degree polynomial samples in three dimensions.

For total-degree polynomials in n variables,

$$\dim \Pi_k(\mathbb{R}^n) = \binom{n+k}{k}.$$

In three dimensions the coefficient counts are:

3.3 A local LSQ-to-Taylor theorem

The practical LSQ construction is already present in [2]. What is not stated there as a standalone theorem is the local convergence statement that makes the LSQ samples mathematically comparable to Taylor samples on regular regions. This section records that result in a form tailored to the later ADSDF discussion.

3.3.1 Setup

Fix a normalized stencil

$$\Xi = \{\xi_1, \dots, \xi_N\} \subset \mathbb{R}^n$$

and positive weights $w_1, \dots, w_N$. Let $W = \text{diag}(w_1, \dots, w_N)$ and let V be the corresponding Vandermonde-type design matrix in the normalized monomial basis. We assume the natural unsolvence condition

$$\text{rank}(V) = \dim \Pi_k(\mathbb{R}^n).$$

Let $U \subset \mathbb{R}^n$ be open, let $K \subset U$ be compact, and assume

$$f \in C^{k+1}(U).$$

For $x \in K$ and $\rho > 0$, define the rescaled function

$$g_{x,\rho}(y) = f(x + \rho y).$$

Let $q_{x,\rho} \in \Pi_k(\mathbb{R}^n)$ be the weighted least-squares fit of $g_{x,\rho}$ on Ξ , and let $\tau_{x,\rho} \in \Pi_k(\mathbb{R}^n)$ be the degree- k Taylor polynomial of $g_{x,\rho}$ at the origin.

Proposition 3.3.1 (Exactness of the normalized LSQ operator).

Under the full-rank assumption on V , the operator

$$L = (V^\top W V)^{-1} V^\top W$$

is bounded and reproduces every polynomial in $\Pi_k(\mathbb{R}^n)$ exactly.

Proof. Boundedness is immediate because all spaces involved are finite-dimensional. Exactness follows from the identity

$$L(V\mathbf{b}) = (V^\top W V)^{-1} V^\top W V \mathbf{b} = \mathbf{b}$$

for every coefficient vector $\mathbf{b}$ in the normalized monomial basis. □

3.3.2 Samplewise comparison

The correct place to compare least-squares samples to Taylor samples is in centered, normalized coordinates. This avoids spurious growth from repeatedly translating coefficients before the estimate has been proved.

Theorem 3.3.1 (LSQ recovers the Taylor polynomial to the expected order).

There exist constants $C > 0$ and $\rho_0 > 0$ such that for all $x \in K$ and all $0 < \rho \leq \rho_0$, the degree- k least-squares fit $q_{x,\rho}$ and the degree- k Taylor polynomial $\tau_{x,\rho}$ satisfy

$$\|q_{x,\rho} - \tau_{x,\rho}\|_{\Pi_k} \leq C\rho^{k+1},$$

where $\|\cdot\|_{\Pi_k}$ is any fixed coefficient norm on $\Pi_k(\mathbb{R}^n)$ in the normalized monomial basis.

Proof. Taylor's theorem applied to $g_{x,\rho}$ at the origin gives

$$g_{x,\rho}(y) = \tau_{x,\rho}(y) + R_{x,\rho}(y), \quad |R_{x,\rho}(y)| \leq C_0\rho^{k+1}$$

uniformly for $x \in K$ and y in bounded subsets of the normalized stencil neighborhood. Sampling on Ξ yields

$$g_{\Xi}(x, \rho) = Vb(x, \rho) + \mathbf{r}(x, \rho), \quad \|\mathbf{r}(x, \rho)\| \leq C_1\rho^{k+1},$$

where $b(x, \rho)$ is the coefficient vector of $\tau_{x,\rho}$. Applying the exact LSQ operator from Proposition 3.3.1,

$$\widehat{b}(x, \rho) = Lg_{\Xi}(x, \rho) = L(Vb(x, \rho) + \mathbf{r}(x, \rho)) = b(x, \rho) + L\mathbf{r}(x, \rho).$$

Hence

$$\|\widehat{b}(x, \rho) - b(x, \rho)\| \leq \|L\| \|\mathbf{r}(x, \rho)\| \leq C\rho^{k+1},$$

which is exactly the desired coefficient estimate. □

The coefficient estimate immediately implies a comparison of the actual polynomials and their derivatives in physical coordinates.

Corollary 3.3.1 (Samplewise polynomial and derivative comparison).

Let

$$p_{x,\rho}(z) = q_{x,\rho}\left(\frac{z-x}{\rho}\right)$$

be the least-squares sample written back in physical coordinates, and let $T_{f;x}^k$ be the degree- k Taylor polynomial of f at x . Then for every multi-index β with $|\beta| \leq k$ there exists $C_\beta > 0$ such that

$$\sup_{z \in x + \rho B_*} |\partial^\beta (p_{x,\rho}(z) - T_{f;x}^k(z))| \leq C_\beta \rho^{k+1-|\beta|},$$

uniformly for $x \in K$ and $0 < \rho \leq \rho_0$, where $B_* \subset \mathbb{R}^n$ is any fixed bounded normalized evaluation set.

Proof. Theorem 3.3.1 says that the normalized coefficients differ by $O(\rho^{k+1})$. Differentiation with respect to the normalized variables changes only the constant, because the evaluation set is bounded. Passing back to physical coordinates contributes the expected factor $\rho^{-|\beta|}$. $\square$

3.3.3 From samples to filtered fields

We now return to ADSDFs. Let $h > 0$ denote the coarse sample spacing, let $\rho = \theta h$ for some fixed $\theta > 0$, and let

$$F_h^T(\mathbf{x}) = \sum_{i \in J_h(\mathbf{x})} \lambda_{i,h}(\mathbf{x}) t_i^h(\mathbf{x}), \quad F_h^L(\mathbf{x}) = \sum_{i \in J_h(\mathbf{x})} \lambda_{i,h}(\mathbf{x}) p_i^h(\mathbf{x})$$

denote the filtered Taylor and LSQ ADSDFs, respectively. Here $J_h(\mathbf{x})$ is the active local footprint, t_i^h is the degree- k Taylor sample at $\mathbf{x}_i$, and p_i^h is the corresponding LSQ sample.

The only assumptions we need on the weights are the usual local ones:

$$\#J_h(\mathbf{x}) \leq J_{\max}, \quad \sum_{i \in J_h(\mathbf{x})} |\lambda_{i,h}(\mathbf{x})| \leq C_\lambda, \quad \sum_{i \in J_h(\mathbf{x})} \|\nabla \lambda_{i,h}(\mathbf{x})\| \leq C_\lambda h^{-1}.$$

These are exactly the bounded-footprint and scaling assumptions one expects from barycentric-type filters.

Finally, we provide a formal statement and proof of the informal LSQ-to-Taylor principle (Theorem 1.4.1) from Chapter 1.

Theorem 3.3.2 (Filtered LSQ fields differ from filtered Taylor fields by the same order).

Let $K \subset U$ be a compact regular region and assume all coarse footprints and fine stencils influencing K remain inside U for sufficiently small h . Then there exist constants $C_0, C_1 > 0$ such that

$$\|F_h^L - F_h^T\|_{L^\infty(K)} \leq C_0 h^{k+1}$$

and

$$\|\nabla F_h^L - \nabla F_h^T\|_{L^\infty(K)} \leq C_1 h^k.$$

Proof. For each active sample, Corollary 3.3.1 gives

$$|p_i^h(\mathbf{x}) - t_i^h(\mathbf{x})| \leq Ch^{k+1}, \quad \|\nabla p_i^h(\mathbf{x}) - \nabla t_i^h(\mathbf{x})\| \leq Ch^k$$

uniformly on the relevant cells. Therefore

$$|F_h^L(\mathbf{x}) - F_h^T(\mathbf{x})| = \left| \sum_{i \in J_h(\mathbf{x})} \lambda_{i,h}(\mathbf{x}) (p_i^h(\mathbf{x}) - t_i^h(\mathbf{x})) \right| \leq \left(\sum_{i \in J_h(\mathbf{x})} |\lambda_{i,h}(\mathbf{x})| \right) Ch^{k+1},$$

which proves the value estimate.

For the gradient, differentiate once:

$$\begin{aligned} \nabla(F_h^L - F_h^T) &= \sum_{i \in J_h(\mathbf{x})} \underbrace{\nabla \lambda_{i,h}(\mathbf{x}) (p_i^h(\mathbf{x}) - t_i^h(\mathbf{x}))}_{\text{weight term}} \\ &\quad + \sum_{i \in J_h(\mathbf{x})} \underbrace{\lambda_{i,h}(\mathbf{x}) (\nabla p_i^h(\mathbf{x}) - \nabla t_i^h(\mathbf{x}))}_{\text{sample term}}. \end{aligned}$$

The assumed bounds on the weights give

$$\|\nabla(F_h^L - F_h^T)\| \leq C_\lambda h^{-1} \cdot Ch^{k+1} + C_\lambda \cdot Ch^k \leq C'h^k.$$

Taking suprema over K yields the claim. $\square$

The upshot is that any local convergence theorem already proved for Taylor ADSDFs transfers immediately to LSQ ADSDFs.

Corollary 3.3.2 (Transfer of Taylor convergence order).

Assume that on a compact regular region K the Taylor ADSDF satisfies

$$\|f - F_h^T\|_{L^\infty(K)} \leq C_T h^{k+1}.$$

Then

$$\|f - F_h^L\|_{L^\infty(K)} \leq (C_T + C_0)h^{k+1}.$$

Likewise, if

$$\|\nabla f - \nabla F_h^T\|_{L^\infty(K)} \leq C'_T h^k,$$

then

$$\|\nabla f - \nabla F_h^L\|_{L^\infty(K)} \leq (C'_T + C_1)h^k.$$

Proof. This is just the triangle inequality applied to Theorem 3.3.2. $\square$

The theorem above is the natural formalization of the statement in [2] that the least-squares fit converges to the Taylor polynomial as the fine stencil shrinks. More importantly for the later chapters, it explains why the practical LSQ construction is a mathematically faithful surrogate for Taylor ADSDFs on regular regions.

3.3.4 Limitations and bridge to the next chapter

Theorem 3.3.2 is intentionally local. It applies only on compact regions where the exact signed distance function is C^{k+1} , so it does not cover cells whose coarse footprint or fine least-squares stencil intersects the cut locus or another nonsmooth feature of the target SDF. Nonconvexity of the target is not itself the issue; the issue is loss of uniqueness of the closest point and therefore loss of SDF smoothness near the cut locus. This is sufficient for the approximation claim made here: on regular regions, the practical LSQ construction has the same local order as the Taylor construction. Near nonsmooth regions the ADSDF may still be useful computationally, but its behavior is empirical rather than covered by this theorem. On compact regular regions:

1. LSQ ADSDFs inherit the approximation order of Taylor ADSDFs;
2. normalized global monomial coordinates are mathematically stable;
3. filtered values, gradients, and Hessians can be efficiently evaluated for geometric queries.

That is why the later energy formulations can use $\hat{\Phi}_{\text{SDF}}$, $\nabla \hat{\Phi}_{\text{SDF}}$, and $\nabla^2 \hat{\Phi}_{\text{SDF}}$ without reverting to finite differences.

Chapter 4

Intrinsic Optimization on Euclidean and Manifold-Valued Variables

This chapter develops the optimization framework used for the retargeting energies introduced in Chapter 2. The search space is not purely Euclidean: vertex positions live in $(\mathbb{R}^3)^N$, but the auxiliary rotational data live in $(S^3)^N$. Consequently, the correct setting is a product of Euclidean and manifold-valued variables.

Two goals drive the chapter. The first is to record the first-order geometry needed to define gradients and descent updates on this product space. The second is to explain why the least-squares structure of the energy favors Gauss–Newton type methods over a full Newton method. Towards that end, we use the existing Riemannian optimization literature, particularly Boumal’s text [17] and related Gauss–Newton literature [18–20], and specialize it to our problems of interest.

4.1 Why manifold optimization appears in retargeting

The state variable of the retargeting problem is

$$\mathcal{X} = (\mathbb{R}^3)^N \times (S^3)^N.$$

A point of $\mathcal{X}$ is written as $(\mathbf{P}, \mathbf{Q})$ with

$$\mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_N), \quad \mathbf{Q} = (q_1, \dots, q_N).$$

The positional factor is Euclidean. The quaternion factor is not: each q_i is constrained to remain on the unit sphere in $\mathbb{H} \cong \mathbb{R}^4$, and hence evolves on a manifold.

A convenient tangent-space representation is the following. A tangent vector at $(\mathbf{P}, \mathbf{Q})$ is written as

$$\eta = (\delta P, \Xi), \quad \delta P = (\delta \mathbf{p}_1, \dots, \delta \mathbf{p}_N) \in (\mathbb{R}^3)^N, \quad \Xi = (X_1, \dots, X_N) \in (\mathbb{R}^3)^N.$$

Here each X_i is identified with a pure imaginary quaternion and represents the tangent vector $q_i X_i \in T_{q_i} S^3$ by left translation. In these coordinates the product metric takes the simple form

$$\langle (\delta P, \Xi), (\delta P', \Xi') \rangle = \sum_{i=1}^N \delta \mathbf{p}_i \cdot \delta \mathbf{p}'_i + \sum_{i=1}^N X_i \cdot X'_i.$$

Throughout the chapter we use this product metric. The energy weights, area weights, and other local coefficients of the energy are absorbed into the residuals themselves rather than into the metric.

With this choice, the gradient splits exactly as expected.

Proposition 4.1.1.

Let $f: \mathcal{X} \rightarrow \mathbb{R}$ be smooth. Then

$$\text{grad } f(\mathbf{P}, \mathbf{Q}) = \left(\text{grad}_{\mathbf{P}} f(\mathbf{P}, \mathbf{Q}), \text{grad}_{\mathbf{Q}} f(\mathbf{P}, \mathbf{Q}) \right),$$

and, more concretely,

$$\text{grad } f(\mathbf{P}, \mathbf{Q}) = \left(\text{grad}_{p_1} f, \dots, \text{grad}_{p_N} f, \text{grad}_{q_1} f, \dots, \text{grad}_{q_N} f \right).$$

Proof. The tangent space of the product splits as

$$T_{(\mathbf{P}, \mathbf{Q})} \mathcal{X} \cong (\mathbb{R}^3)^N \times \prod_{i=1}^N T_{q_i} S^3,$$

and the product metric is the sum of the constituent metrics. Writing out the defining identity for $\text{grad } f$ against an arbitrary tangent vector therefore shows that the Euclidean and quaternion components decouple exactly as in ordinary multivariable calculus. $\square$

Historically, one may motivate manifold algorithms through local parameterizations: near a point $x \in \mathcal{M}$, replace the manifold by local coordinates, perform a Euclidean step, and map back. That point of view is useful for intuition, and it underlies classical Newton-like methods on manifolds such as those of Hüper and Trumpf [8]. For our work, however, the computationally natural language is that of *retractions* and *Riemannian gradients*, following Boumal’s treatment [17, Chs. 3–4]. This avoids carrying local coordinates explicitly while still retaining an intrinsic update rule.

4.2 Riemannian gradients, retractions, and descent

In Euclidean space, gradient descent updates

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k \nabla f(\mathbf{x}_k)$$

by subtracting a multiple of the gradient. On a manifold there is no globally meaningful subtraction of points. What survives is the local picture: the gradient is a tangent vector, and a *retraction* maps a tangent step back to the manifold.

Definition 4.2.1.

A *Riemannian manifold* is a smooth manifold $\mathcal{M}$ equipped with an inner product $\langle \cdot, \cdot \rangle_x$ on each tangent space $T_x \mathcal{M}$, varying smoothly with x . For a smooth function $f: \mathcal{M} \rightarrow \mathbb{R}$, the *Riemannian gradient* $\text{grad } f(x) \in T_x \mathcal{M}$ is the unique tangent vector satisfying

$$Df(x)[\xi] = \langle \text{grad } f(x), \xi \rangle_x \quad \text{for all } \xi \in T_x \mathcal{M}.$$

A *retraction* on $\mathcal{M}$ is a smooth map

$$R: T\mathcal{M} \rightarrow \mathcal{M}, \quad R_x(0) = x, \quad DR_x(0) = \text{Id}_{T_x\mathcal{M}}.$$

The Riemannian gradient is the manifold analogue of the usual Euclidean gradient, and the retraction replaces the Euclidean update $x + s$ by a geometrically compatible map $R_x(s)$.

The Riemannian exponential $\text{Exp}_x(v)$ is the endpoint at time 1 of the geodesic starting at x with initial velocity $v \in T_x\mathcal{M}$, whenever that geodesic is defined. A retraction is a first-order approximation to this map. We use Exp for the Riemannian exponential on a manifold and $\exp$ for the Lie-group or quaternion exponential at the identity.

For the product space $\mathcal{X} = (\mathbb{R}^3)^N \times (S^3)^N$, the componentwise retraction we use is

$$R_{(\mathbf{P}, \mathbf{Q})}(\delta P, \Xi) = \left(\mathbf{P} + \delta P, (q_1 \exp(X_1), \dots, q_N \exp(X_N)) \right).$$

Thus the Euclidean variables are updated by addition, while the quaternion variables are updated by the quaternion exponential. If

$$\text{grad } \mathcal{E}_{\text{total}}(\mathbf{P}, \mathbf{Q}) = (\nabla_{\mathbf{P}} \mathcal{E}_{\text{total}}(\mathbf{P}, \mathbf{Q}), (q_1 X_1^{\text{tot}}, \dots, q_N X_N^{\text{tot}})),$$

then one step of Riemannian gradient descent takes the explicit form

$$\mathbf{p}_i^+ = \mathbf{p}_i - \alpha \nabla_{\mathbf{p}_i} \mathcal{E}_{\text{total}}, \quad q_i^+ = q_i \exp(-\alpha X_i^{\text{tot}}).$$

Only the Lie algebra representatives $X_i^{\text{tot}} \in \mathbb{R}^3$ are needed for the quaternion update.

Theorem 4.2.1 (Boumal, restated in the present notation).

Let $\mathcal{M}$ be a Riemannian manifold, let $f: \mathcal{M} \rightarrow \mathbb{R}$ be smooth and bounded below, and let R

be a retraction on $\mathcal{M}$. Assume there exists $L > 0$ such that the pullbacks of f satisfy

$$f(R_x(s)) \leq f(x) + \langle \text{grad } f(x), s \rangle_x + \frac{L}{2} \|s\|^2$$

for all iterates x and all tangent steps s under consideration. Then the constant-step iteration

$$x_{k+1} = R_{x_k} \left(-\frac{1}{L} \text{grad } f(x_k) \right)$$

satisfies

$$\min_{0 \leq j < K} \|\text{grad } f(x_j)\| \leq \sqrt{\frac{2L(f(x_0) - f_{\text{low}})}{K}},$$

and in particular $\|\text{grad } f(x_k)\| \rightarrow 0$. Every accumulation point of the iterates is therefore a critical point of f .

Proof. This is a direct restatement of Boumal’s sufficient-decrease analysis for Riemannian gradient descent; see [17, Proposition 4.7 and Corollary 4.9]. $\square$

The theorem is the first-order convergence result we need. A backtracking line-search variant follows under the same pullback-Lipschitz hypothesis; see [17, Corollary 4.13]. Stronger local statements exist as well. For example, if a critical point has positive-definite Riemannian Hessian, then constant-step Riemannian gradient descent converges locally at least linearly [17, Theorem 4.20]. We do not make that kind of statement central here because the retargeting objective is nonconvex and later sections will lean instead on the least-squares structure of the energy.

4.3 Lie groups with bi-invariant metrics

The manifold S^3 is not only a Riemannian manifold but also a Lie group. This is what makes the quaternion factor computationally pleasant.

Let G be a Lie group with identity element e , group multiplication written multiplicatively, and left translation

$$L_p: G \rightarrow G, \quad L_p(r) = pr.$$

The tangent space $T_e G$ is the Lie algebra of G . With respect to a bi-invariant Riemannian metric on G , left translation is an isometry, so tangent vectors at arbitrary points can be represented by vectors at the identity. In this situation the Riemannian exponential and the Lie-group exponential are compatible in the following sense:

Lemma 4.3.1.

Let G be a Lie group with a bi-invariant Riemannian metric. If $V_p \in T_p G$ has Lie algebra representative $V_e = [dL_{p^{-1}}]_p(V_p) \in T_e G$, then

$$\text{Exp}_p(V_p) = p \exp(V_e).$$

The metric is what allows one to define gradients and invoke convergence results such as Theorem 4.2.1. The group structure is what makes the update map explicit. Bi-invariance ensures that these two structures cooperate cleanly.

4.4 Working concretely in S^3

We now specialize to S^3 , viewed as the Lie group of unit quaternions. Vectors in $\mathbb{R}^3$ are identified with pure imaginary quaternions. Thus, if $v \in \mathbb{R}^3$ and $q \in S^3$, the rotated vector is

$$v \longmapsto qvq^*.$$

The tangent space at the identity is the three-dimensional space of pure imaginary quaternions, which we identify with $\mathbb{R}^3$ itself. If $X \in \mathbb{R}^3 \cong T_e S^3$, then the corresponding tangent vector at q is $qX \in T_q S^3$.

The exponential map at a point therefore becomes

$$\text{Exp}_q(qX) = q \exp(X).$$




Figure 4.1: A unit quaternion visualized by its action on a coordinate axis.

For the unit quaternions this is just quaternion arithmetic. This is the practical reason for introducing the Lie-group viewpoint.

A separate convention is needed for the logarithm. For a unit quaternion $r \in S^3$, write $r = s + \mathbf{v}$, where $s \in \mathbb{R}$ is the scalar part and $\mathbf{v} \in \mathbb{R}^3$ is the vector part. If $\mathbf{v} \neq 0$, write $\mathbf{u} = \mathbf{v}/\|\mathbf{v}\|$ and

$$r = \cos \theta + \mathbf{u} \sin \theta, \quad \theta = \text{atan2}(\|\mathbf{v}\|, s).$$

The principal logarithm is $\log r = \theta \mathbf{u}$. If $\mathbf{v} = 0$, define $\log(1) = 0$; the point $r = -1$ is the antipodal branch point and must be handled by the chosen branch convention. When $r = q_i^* q_j$ represents a relative rotation, the signs of the quaternions should be chosen consistently so that the scalar part of r is nonnegative before applying the logarithm. This selects the shortest branch. With that convention, $\|\log(q_i^* q_j)\|$ is the geodesic distance on S^3 between q_i and q_j , and twice that norm is the relative rotation angle in $\text{SO}(3)$.

4.5 Quaternion identities for gradients and Jacobians

The energy terms from Chapter 2 repeatedly use expressions of the form qvq^* , so it is worth recording a few reusable formulas.

Lemma 4.5.1.

Fix $v \in \mathbb{R}^3$ and define

$$a(q) = qvq^*, \quad q \in S^3.$$

Let $X \in T_e S^3 \cong \mathbb{R}^3$ and consider the tangent vector $qX \in T_q S^3$. Then

$$Da(q)[qX] = 2q(X \times v)q^*.$$

Proof. Write $q(t) = q \exp(tX)$. Using the first-order expansion of the quaternion exponential,

$$\exp(tX)v \exp(-tX) = \underbrace{v}_{\text{value at } t=0} + \underbrace{t(Xv - vX)}_{= 2t(X \times v)} + O(t^2).$$

Multiplying on the left by q and on the right by q^* yields

$$q(t)vq(t)^* = qvq^* + 2t q(X \times v)q^* + O(t^2),$$

which is the claimed differential formula. □

Proposition 4.5.1.

Let $u, v \in \mathbb{R}^3$ be fixed and define

$$f(q) = (qvq^*) \cdot u.$$

Then the Lie algebra representative of $\text{grad } f(q)$ is

$$X_f(q) = 2(v \times q^* u q).$$

Proof. For a tangential variation qX , Lemma 4.5.1 gives

$$Df(q)[qX] = 2q(X \times v)q^* \cdot u = 2(X \times v) \cdot (q^* u q) = 2X \cdot (v \times q^* u q).$$

Because left translation is an isometry for the bi-invariant metric on S^3 , the inner product at q is represented at the identity by the ordinary dot product on $\mathbb{R}^3$. Hence the displayed vector is precisely the Lie algebra representative of the Riemannian gradient. $\square$

Corollary 4.5.1.

Fix $u, v \in \mathbb{R}^3$ and define

$$g(q) = \frac{1}{2} \|qvq^* - u\|^2.$$

Then the Lie algebra representative of $\text{grad } g(q)$ is

$$X_g(q) = -2(v \times q^* u q).$$

Proof. Expand the square:

$$\frac{1}{2} \|qvq^* - u\|^2 = \underbrace{\frac{1}{2} \|qvq^*\|^2}_{\text{constant in } q} + \underbrace{\frac{1}{2} \|u\|^2}_{\text{constant}} - (qvq^*) \cdot u.$$

The first term is constant because quaternion conjugation preserves Euclidean length. The claim now follows from Proposition 4.5.1. $\square$

Proposition 4.5.2.

Fix quaternions $q_1, \dots, q_m \in S^3$ contained in a common geodesically convex neighborhood,

unit-quaternion frames: red e_x , green e_y , blue e_z




Figure 4.2: The Karcher mean of two unit quaternions q_1, q_2 .

and let

$$\mathcal{E}_K(q) = \frac{1}{2m} \sum_{j=1}^m \|\log(q^* q_j)\|^2.$$

Then the Riemannian gradient has Lie algebra representative

$$X_{\mathcal{E}_K}(q) = -\frac{1}{m} \sum_{j=1}^m \log(q^* q_j).$$

Consequently, a weighted Karcher mean $\bar{q}$ satisfies the first-order condition

$$\sum_{j=1}^m \alpha_j \log(\bar{q}^* q_j) = 0, \quad \alpha_j \geq 0, \quad \sum_j \alpha_j = 1.$$

Proof. On any Riemannian manifold, the gradient of one half the squared distance to a fixed point is minus the Riemannian logarithm, away from the cut locus; see, for example, [17, Sec. 10.2] and the classical treatment of Karcher means [7]. On S^3 with the bi-invariant metric, the Riemannian logarithm at q is

$$\text{Log}_q(q_j) = q \log(q^* q_j).$$

Therefore the tangent gradient of each summand is $-q \log(q^* q_j)$, and its Lie algebra representative is $-\log(q^* q_j)$. Averaging the summands gives the stated formula. $\square$

The preceding proposition is the only worked example we need. It is simple enough to illustrate intrinsic averaging, and it reappears later in the manifold-valued subdivision chapter.

4.6 Gradients of the retargeting energies

We now differentiate the energy terms from Chapter 2. Because Chapter 3 has already replaced the ideal smooth target by a discrete algebraic field, all target-side quantities are now written in terms of $\hat{\Phi}_{\text{SDF}}$, $\nabla \hat{\Phi}_{\text{SDF}}$, and $\nabla^2 \hat{\Phi}_{\text{SDF}}$.

To keep the formulas short, write

$$\phi_i = \hat{\Phi}_{\text{SDF}}(\mathbf{p}_i), \quad g_i = \nabla \hat{\Phi}_{\text{SDF}}(\mathbf{p}_i), \quad \mathbf{H}_i = \nabla^2 \hat{\Phi}_{\text{SDF}}(\mathbf{p}_i), \quad a_i = \text{area}(\star i).$$

When the gradient of the discrete field is normalized in an implementation, the formulas below acquire an extra chain-rule factor. The unnormalized form already captures the geometry cleanly and is the one used here.

4.6.1 Distance, clearance, and landmarks

For derivative calculations it is convenient to write the least-squares terms with a factor of $\frac{1}{2}$.

Distance / offset. Let

$$\mathcal{E}_{\text{dist}}(\mathbf{P}) = \frac{1}{2} \sum_{i \in \mathcal{V}} \underbrace{(\phi_i - d_i)^2}_{\text{distance to the desired level set}} a_i.$$

Then

$$\nabla_{\mathbf{p}_i} \mathcal{E}_{\text{dist}} = \underbrace{(\phi_i - d_i)}_{\text{scalar residual}} \underbrace{g_i}_{\text{target gradient}} a_i.$$

Clearance. The one-sided penalty from Chapter 2 is not smooth. To keep the least-squares block compatible with Gauss–Newton, we replace the hard hinge by a smooth residual surrogate.

A convenient choice is the softplus

$$\rho_\varepsilon(t) = \varepsilon \log(1 + e^{t/\varepsilon}), \quad \rho'_\varepsilon(t) = \frac{1}{1 + e^{-t/\varepsilon}},$$

and the corresponding clearance energy

$$\mathcal{E}_{\text{clear},\varepsilon}(\mathbf{P}) = \frac{1}{2} \sum_{i \in \mathcal{V}} \rho_\varepsilon(d_i - \phi_i)^2 a_i.$$

Then

$$\nabla_{\mathbf{p}_i} \mathcal{E}_{\text{clear},\varepsilon} = -\rho_\varepsilon(d_i - \phi_i) \rho'_\varepsilon(d_i - \phi_i) g_i a_i.$$

This redesign is a modeling choice rather than the canonical barrier formulation. It is adopted here only to keep the main algorithmic development focused on a smooth least-squares block; more general barrier and smooth-composite variants are possible, but they are not needed for the main line of the thesis. One could keep the hard hinge and solve a semismooth least-squares problem. The softplus version is used here only to keep the main solver within a smooth Gauss–Newton or Levenberg–Marquardt framework; the tradeoff is that the active-set transition is smoothed over a small width controlled by ε .

Landmarks. For a point-based landmark term

$$\mathcal{E}_{\text{land}}(\mathbf{P}) = \frac{1}{2} \sum_{\ell=1}^m \omega_\ell \|\mathbf{p}_{i_\ell} - \mathbf{x}_\ell^*\|^2,$$

we have the obvious Euclidean gradient

$$\nabla_{\mathbf{p}_{i_\ell}} \mathcal{E}_{\text{land}} = \omega_\ell (\mathbf{p}_{i_\ell} - \mathbf{x}_\ell^*).$$

If several landmarks share the same vertex, their contributions simply add.

4.6.2 Normal alignment

For the normal term it is convenient to use the vector-residual form

$$\mathcal{E}_{\text{norm}}(\mathbf{P}, \mathbf{Q}) = \frac{1}{2} \sum_{i \in \mathcal{V}} \chi_i \underbrace{\|q_i \mathbf{n}_i q_i^* - g_i\|^2}_{\text{rotated source normal minus target normal proxy}} a_i.$$

On an exact smooth SDF, where $\|g_i\| = 1$, this is equivalent to the scalar form $1 - (q_i \mathbf{n}_i q_i^*) \cdot g_i$ up to an additive constant. In the discrete setting the vector-residual form is more convenient for Gauss–Newton.

Define the residual

$$\vec{\mathbf{r}}_i^{\text{norm}} = q_i \mathbf{n}_i q_i^* - g_i.$$

Then

$$\nabla_{\mathbf{p}_i} \mathcal{E}_{\text{norm}} = -\chi_i a_i \mathbf{H}_i \vec{\mathbf{r}}_i^{\text{norm}},$$

using symmetry of the Hessian. For the quaternion factor, Corollary 4.5.1 gives the Lie algebra representative

$$X_i^{\text{norm}} = -2\chi_i a_i (\mathbf{n}_i \times q_i^* g_i q_i).$$

Thus the actual tangent gradient at q_i is $q_i X_i^{\text{norm}}$.

4.6.3 The ARAP term

Write the reference and deformed edges as

$$\vec{\mathbf{e}}_{ij}^{(0)} = \mathbf{p}_i^{(0)} - \mathbf{p}_j^{(0)}, \quad \vec{\mathbf{e}}_{ij} = \mathbf{p}_i - \mathbf{p}_j, \quad \alpha_{ij} = \mathbf{w}_{ij} a_i.$$

The near-isometric term becomes

$$\mathcal{E}_{\text{ARAP}}(\mathbf{P}, \mathbf{Q}) = \frac{1}{2} \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{N}_1(i)} \alpha_{ij} \left\| \underbrace{q_i \vec{\mathbf{e}}_{ij}^{(0)} q_i^*}_{\text{rotated reference edge}} - \underbrace{\vec{\mathbf{e}}_{ij}}_{\text{deformed edge}} \right\|^2.$$

Define the oriented residual

$$\vec{\mathbf{r}}_{ij}^{\text{ARAP}} = q_i \vec{\mathbf{e}}_{ij}^{(0)} q_i^* - \vec{\mathbf{e}}_{ij}.$$

Its differential along a tangent direction $(\delta \mathbf{p}_i, \delta \mathbf{p}_j, q_i X_i)$ is, by Lemma 4.5.1,

$$D\vec{\mathbf{r}}_{ij}^{\text{ARAP}}[\delta \mathbf{p}_i, \delta \mathbf{p}_j, q_i X_i] = 2q_i(X_i \times \vec{\mathbf{e}}_{ij}^{(0)})q_i^* - (\delta \mathbf{p}_i - \delta \mathbf{p}_j).$$

From this one reads off the local contributions to the gradient. The quaternion representative at the anchor vertex i is

$$X_i^{\text{ARAP}} = -2 \sum_{j \in \mathcal{N}_1(i)} \alpha_{ij} \left(\vec{\mathbf{e}}_{ij}^{(0)} \times q_i^* \vec{\mathbf{e}}_{ij} q_i \right).$$

For the positional factor, the oriented pair (i, j) contributes

$$\nabla_{\mathbf{p}_i} \left(\frac{1}{2} \alpha_{ij} \|\vec{\mathbf{r}}_{ij}^{\text{ARAP}}\|^2 \right) = -\alpha_{ij} \vec{\mathbf{r}}_{ij}^{\text{ARAP}}, \quad \nabla_{\mathbf{p}_j} \left(\frac{1}{2} \alpha_{ij} \|\vec{\mathbf{r}}_{ij}^{\text{ARAP}}\|^2 \right) = \alpha_{ij} \vec{\mathbf{r}}_{ij}^{\text{ARAP}}.$$

Hence the full gradient at a vertex is assembled by summing its own anchor contributions together with the neighboring terms in which the vertex appears as the second endpoint.

4.6.4 Smoothing and Karcher-type terms

A typical smoothing term is built from geodesic discrepancies between neighboring quaternions. For a single pair (q_i, q_j) , define

$$\psi_{ij}(q_i) = \frac{1}{2} \|\log(q_i^* q_j)\|^2.$$

Then Proposition 4.5.2 yields the Lie algebra representative

$$X_{ij}^{\text{smooth}} = -\log(q_i^* q_j).$$

A smoothing energy is obtained by summing such local contributions over neighboring pairs. Likewise, a weighted Karcher mean problem is nothing more than the minimization of a weighted sum of these pairwise half-squared log terms.

4.7 Karcher means as the recurring worked example

The discussion above already explains why Karcher means are the right intrinsic replacement for Euclidean averaging. Given quaternions $q_1, \dots, q_m$ with weights $\alpha_1, \dots, \alpha_m$, the weighted Karcher mean minimizes

$$q \mapsto \frac{1}{2} \sum_{j=1}^m \alpha_j \|\log(q^* q_j)\|^2.$$

By Proposition 4.5.2, its first-order condition is

$$\sum_{j=1}^m \alpha_j \log(q^* q_j) = 0.$$

This makes Karcher means an ideal intermediary example for two reasons. First, they demonstrate intrinsic optimization on a genuinely curved variable space without introducing the full retargeting energy. Second, the same averaging principle reappears later when subdivision rules for Euclidean control meshes are lifted to manifold-valued auxiliary data.

4.8 Why Gauss–Newton rather than Newton?

Most of the retargeting energies from Chapter 2 are least-squares in character. This is the first reason to prefer Gauss–Newton: it is built for objectives of the form

$$f(x) = \frac{1}{2} \|F(x)\|^2, \quad x \in \mathcal{X},$$

where F is a residual map with values in a Euclidean space.

Residual	Gauss–Newton needs	Full Newton needs
$\widehat{\Phi}_{\text{SDF}}(\mathbf{p}) - d$	$\nabla \widehat{\Phi}_{\text{SDF}}$	$\nabla^2 \widehat{\Phi}_{\text{SDF}}$
$q\mathbf{n}q^* - \nabla \widehat{\Phi}_{\text{SDF}}(\mathbf{p})$	$\nabla^2 \widehat{\Phi}_{\text{SDF}}$	3rd derivatives of $\widehat{\Phi}_{\text{SDF}}$
curvature term using $\nabla^2 \widehat{\Phi}_{\text{SDF}}$	3rd derivatives of $\widehat{\Phi}_{\text{SDF}}$	4th derivatives of $\widehat{\Phi}_{\text{SDF}}$

Table 4.1: Derivative order required by Gauss–Newton versus full Newton for the most important residuals in this thesis. The derivative-count argument is one of the main reasons to favor Gauss–Newton in an ADSDF-based formulation.

A full Riemannian Newton step would instead solve

$$\text{Hess } f(x)[\eta] = -\text{grad } f(x)$$

in the tangent space and then retract. For a least-squares objective, the Hessian of this full Newton method decomposes schematically as

$$\text{Hess } f(x) = \underbrace{DF(x)^* DF(x)}_{\text{Gauss–Newton part}} + \underbrace{\sum_a F_a(x) \text{ Hess } F_a(x)}_{\text{second derivatives of the residuals}} .$$

Gauss–Newton keeps only the first term. A zero-residual solution is a point x_* where every residual component vanishes, i.e. $F(x_*) = 0$. At such a point the second term is exactly zero. Away from a zero-residual solution it is generally not zero; it becomes small only when the residual components themselves are small. This is why Gauss–Newton has Newton-like local behavior asymptotically near an exact fit, while away from zero residual it remains a structured approximation rather than the full Newton method [18, 19].

For this thesis, the decisive point is that Gauss–Newton saves one derivative order everywhere. The most relevant residuals are summarized in Table 4.1.

This derivative-count argument should be read together with the practical conclusions of the ADSDF chapter. Valasek and Bán find first- and second-order fields to be the practical sweet spot, while third order is already significantly heavier in storage and evaluation cost [2]. A formulation

whose natural second-order method immediately demands one more derivative order is therefore poorly aligned with the representation strategy of this thesis.

4.9 Riemannian Gauss–Newton and Levenberg–Marquardt

Let $\mathcal{M}$ be a Riemannian manifold, let $F: \mathcal{M} \rightarrow \mathbb{R}^m$ be a smooth residual map, and define

$$f(x) = \frac{1}{2} \|F(x)\|^2.$$

At a point $x \in \mathcal{M}$, Gauss–Newton linearizes the residual in the tangent space and chooses the tangent step

$$\eta_{\text{GN}}(x) \in \arg \min_{\eta \in T_x \mathcal{M}} \frac{1}{2} \|F(x) + DF(x)[\eta]\|^2.$$

The next iterate is then obtained by retraction:

$$x^+ = R_x(\eta_{\text{GN}}).$$

When the differential $DF(x)$ has full column rank, this is equivalent to the normal equations

$$DF(x)^* DF(x)[\eta] = -DF(x)^* F(x).$$

Here $DF(x)^*$ denotes the adjoint of $DF(x)$ with respect to the chosen Riemannian metric on $T_x \mathcal{M}$ and the Euclidean inner product on the residual space. A damped or Levenberg–Marquardt step inserts a positive parameter μ :

$$(DF(x)^* DF(x) + \mu I)[\eta] = -DF(x)^* F(x).$$

The damping parameter μ may be viewed either as Tikhonov regularization of the tangent-space least-squares problem or as a trust-region-like safeguard.

Theorem 4.9.1 (Breiding–Vannieuwenhoven [19], simplified).

Let $\mathcal{M}$ be an embedded Riemannian submanifold and let $f(x) = \frac{1}{2}\|F(x)\|^2$ with $F: \mathcal{M} \rightarrow \mathbb{R}^m$ smooth. Suppose $x_\star$ is a local minimizer at which $DF(x_\star)$ is injective. Then the Riemannian Gauss–Newton iteration is locally linearly convergent if $F(x_\star) \neq 0$, and locally quadratically convergent if $F(x_\star) = 0$.

Proof. This is the qualitative content of the local convergence theorem proved in [19]. The cited paper gives explicit rate and radius constants in terms of the conditioning of the differential. $\square$

Exact contact or exact landmark satisfaction corresponds to a zero-residual regime, where Gauss–Newton has Newton-like local behavior asymptotically. More generally, one should only expect linear local behavior.

4.10 Residual formulation of the retargeting objective

With the preceding discussion in place, the least-squares block of the retargeting objective can be written explicitly. Up to the area factors, energy weights, and the smoothing choice described above, the residuals are

$$\begin{aligned}
r_i^{\text{dist}}(\mathbf{P}) &= \underbrace{\sqrt{T_{\text{dist}} a_i}}_{\text{weight}} \underbrace{(\widehat{\Phi}_{\text{SDF}}(\mathbf{p}_i) - d_i)}_{\text{scalar distance residual}}, \\
r_i^{\text{clear}}(\mathbf{P}) &= \sqrt{T_{\text{clear}} a_i} \underbrace{\rho_\varepsilon(d_i - \widehat{\Phi}_{\text{SDF}}(\mathbf{p}_i))}_{\text{smooth one-sided clearance residual}}, \\
\vec{r}_i^{\text{norm}}(\mathbf{P}, \mathbf{Q}) &= \sqrt{T_{\text{norm}} \chi_i a_i} \underbrace{\left(q_i \mathbf{n}_i q_i^* - \nabla \widehat{\Phi}_{\text{SDF}}(\mathbf{p}_i) \right)}_{\text{rotated source normal minus target normal}}, \\
\vec{r}_{ij}^{\text{ARAP}}(\mathbf{P}, \mathbf{Q}) &= \sqrt{T_{\text{ARAP}} \mathbf{w}_{ij} a_i} \underbrace{\left(q_i \vec{e}_{ij}^{(0)} q_i^* - \vec{e}_{ij} \right)}_{\text{rotated reference edge minus deformed edge}}, \\
\vec{r}_\ell^{\text{land}}(\mathbf{P}) &= \sqrt{T_{\text{land}} \omega_\ell} \underbrace{(\mathbf{p}_{i_\ell} - \mathbf{x}_\ell^*)}_{\text{landmark misfit}}.
\end{aligned}$$

Collecting these into one residual map $F(\mathbf{P}, \mathbf{Q})$, the objective becomes

$$\mathcal{E}_{\text{LS}}(\mathbf{P}, \mathbf{Q}) = \frac{1}{2} \|F(\mathbf{P}, \mathbf{Q})\|^2.$$

Optional smoothing terms based on $\log(q_i^* q_j)$ may be appended as additional residual blocks as well, but their Jacobians require the differential of the logarithm map. To keep the notation focused on the main fitting mechanism, we leave that extension outside the core Gauss–Newton block. The design choice has been to rewrite the important one-sided terms as smooth residual penalties so that the central solver remains a damped Gauss–Newton method rather than a hybrid Gauss–Newton / barrier method.

Computing the Jacobian blocks from the above residuals is straightforward. Using the notation introduced above:

$$\begin{aligned} r_i^{\text{dist}}(\mathbf{p}_i + \delta \mathbf{p}_i) &= r_i^{\text{dist}} + \sqrt{T_{\text{dist}} a_i} g_i \cdot \delta \mathbf{p}_i + O(\|\delta \mathbf{p}_i\|^2), \\ r_i^{\text{clear}}(\mathbf{p}_i + \delta \mathbf{p}_i) &= r_i^{\text{clear}} - \sqrt{T_{\text{clear}} a_i} \rho'_\varepsilon(d_i - \phi_i) g_i \cdot \delta \mathbf{p}_i + O(\|\delta \mathbf{p}_i\|^2), \\ \vec{r}_i^{\text{norm}}(\mathbf{p}_i + \delta \mathbf{p}_i, q_i \exp(X_i)) &= \vec{r}_i^{\text{norm}} + \sqrt{T_{\text{norm}} \chi_i a_i} (2q_i(X_i \times \mathbf{n}_i) q_i^* \\ &\quad - \mathbf{H}_i \delta \mathbf{p}_i) + O(\|(\delta \mathbf{p}_i, X_i)\|^2), \\ \vec{r}_{ij}^{\text{ARAP}}(\mathbf{P} + \delta \mathbf{P}, q_i \exp(X_i)) &= \vec{r}_{ij}^{\text{ARAP}} + \sqrt{T_{\text{ARAP}} \mathbf{w}_{ij} a_i} (2q_i(X_i \times \vec{\mathbf{e}}_{ij}^{(0)}) q_i^* \\ &\quad - (\delta \mathbf{p}_i - \delta \mathbf{p}_j)) + O(\|(\delta \mathbf{p}_i, \delta \mathbf{p}_j, X_i)\|^2), \\ \vec{r}_\ell^{\text{land}}(\mathbf{P} + \delta \mathbf{P}) &= \vec{r}_\ell^{\text{land}} + \sqrt{T_{\text{land}} \omega_\ell} \delta \mathbf{p}_{i_\ell} + O(\|\delta \mathbf{P}\|^2). \end{aligned}$$

These formulas are exactly the tangent-space linearizations that populate the Jacobian $DF(\mathbf{P}, \mathbf{Q})$ and therefore the normal equations.

The corresponding tangent step $(\delta \mathbf{P}, \Xi) \in T_{(\mathbf{P}, \mathbf{Q})} \mathcal{X}$ is obtained from the damped normal equations

$$(DF(\mathbf{P}, \mathbf{Q})^* DF(\mathbf{P}, \mathbf{Q}) + \mu I) \begin{bmatrix} \delta \mathbf{P} \\ \Xi \end{bmatrix} = -DF(\mathbf{P}, \mathbf{Q})^* F(\mathbf{P}, \mathbf{Q}),$$

followed by the product retraction

$$\mathbf{p}_i^+ = \mathbf{p}_i + \delta \mathbf{p}_i, \quad q_i^+ = q_i \exp(X_i).$$

The residual differentials recorded earlier show that $DF(\mathbf{P}, \mathbf{Q})$ has a sparse structure induced by one-ring mesh couplings and local SDF queries. That sparsity is one of the main reasons the Gauss–Newton viewpoint is practical for this problem.

4.11 What this chapter supplies for the rest of the thesis

This chapter has supplied exactly the optimization ingredients needed later on:

- a Riemannian/product-manifold framework for the mixed variable space $(\mathbb{R}^3)^N \times (S^3)^N$,
- explicit quaternion identities that make gradient computations concrete,
- closed-form first-order derivatives for the core energy terms,
- a clear reason to prefer Gauss–Newton over full Newton in an ADSDF-based setting, and
- a damped Gauss–Newton / Levenberg–Marquardt update written directly on the product space.

The later computational chapters can therefore focus on implementation choices and experiments rather than reintroducing the geometry.

Chapter 5

Subdivision and Multiresolution Refinement of Manifold-Valued Data

This chapter is optional on a first reading. The central retargeting method of the thesis does not depend on it directly, and nothing in the later chapters requires the proof given here in order to be read. Its role is to be a companion to the main material: it explains how the manifold-valued data introduced earlier can be refined coherently in a multiresolution pipeline, and it records a convergence result showing that this refinement has a continuous limit under the stated hypotheses.

The motivating question is simple. Once the conforming mesh carries both Euclidean vertex positions and manifold-valued auxiliary data, it is natural to solve a coarse fitting problem first and only later refine the mesh if more detail is needed. Classical subdivision already provides such a coarse-to-fine mechanism for Euclidean control meshes. The present chapter asks for the corresponding statement when the refined data live on a Riemannian manifold and Euclidean affine averages must be replaced by Karcher means.

We proceed in three stages. We begin with a visual introduction to subdivision surfaces, with emphasis on the Loop and Catmull–Clark schemes. We then isolate the small amount of classical linear theory needed later: regular versus extraordinary vertices, local subdivision matrices, and the characteristic-map viewpoint. Finally, we prove a continuous convergence theorem for a manifold-valued analogue obtained by replacing positive affine averages with Karcher means. The proof is organized to make clear which ingredients are imported from the classical linear theory and which are established here.

5.1 Why subdivision appears in this thesis

We motivate subdivision as a crude multiresolution technique. Suppose a retargeting problem has been solved on a coarse mesh, producing a set of deformed vertex positions together with




Figure 5.1: (Left) A conforming surface is to be “shrinkwrapped” to a target monkey head surface. (Center) The procedure begins with a coarse resolution where the conforming surface cannot be cleanly fit to the target. (Right) The conforming surface is refined and is able to be cleanly shrinkwrapped to the target.

quaternion data that encode local rotational information. If one wishes to continue the optimization at a higher resolution, both types of data must be refined in a compatible way. Refining only the positions would destroy the coupling between geometry and auxiliary manifold-valued data that the earlier energy design relies on.

Classical subdivision offers the right Euclidean template. A finite control mesh is repeatedly refined by local weighted averages, and the resulting sequence converges to a smooth limit surface under the usual hypotheses [9–11, 21, 22]. The manifold-valued analogue considered here keeps the same combinatorics and the same positive weights, but replaces affine averages by Karcher means. The resulting refinement rule is nonlinear, yet it remains close to the original linear scheme when the control mesh is already fine. The proof below makes that idea precise.

5.2 A visual introduction to subdivision surfaces

5.2.1 Repeated refinement and limit surfaces

Subdivision defines a curve or surface as the limit of successive refinements of a control mesh. At each step, new vertices are created and old ones may be repositioned according to a fixed local rule. Repeating the same rule produces a sequence of finer and finer meshes which, under appropriate hypotheses, converges to a continuous or smooth limit surface [11, 23].

For the purposes of this chapter, the key points are visual and local:

- subdivision is based on *stencils*, that is, short weighted averaging rules applied repeatedly;
- the same topological refinement pattern is used at every level;




Figure 5.2: A mesh $\mathcal{M}^0$ is successively refined into $\mathcal{M}^1, \mathcal{M}^2, \dots$. In the limit, it becomes a smooth surface $\mathcal{M}^\infty$.

- the surface is not given in closed form at the outset, but emerges as the limit of refinement.

This viewpoint is especially well suited to arbitrary topology, because the same local averaging language can be adapted to neighborhoods that are not part of a regular tensor-product grid.

5.2.2 Loop and Catmull–Clark

The two schemes we keep in mind throughout are Loop subdivision for triangular meshes and Catmull–Clark subdivision for quadrilateral meshes. Loop subdivision inserts edge vertices and replaces each triangle by four smaller triangles. On regular triangular regions it behaves like a box-spline-type scheme, and it is the standard triangular approximating subdivision method in geometry processing and graphics [10, 11]. Catmull–Clark performs the analogous role for quadrilateral meshes: it inserts face points and edge points, updates vertex positions, and reproduces bicubic tensor-product spline behavior on regular quadrilateral regions [9, 11].

This chapter does not require their full masks in detail. The important point is that both are:

- stationary,
- affine-invariant,
- finitely supported,
- local weighted-averaging schemes.

These are exactly the structural properties that carry over to the manifold-valued setting.




Figure 5.3: Splitting a quadrilateral into four smaller quadrilaterals for Catmull–Clark subdivision. A point has been inserted at the center of the faces and at the midpoint of the 4 edges.

5.2.3 Regular and extraordinary vertices

The local behavior of a subdivision rule is simplest on a regular mesh patch. For triangular schemes such as Loop, a regular interior vertex has valence 6. For quadrilateral schemes such as Catmull–Clark, a regular interior vertex has valence 4. Any vertex of different valence is called *extraordinary*. The neighborhood of an extraordinary vertex is precisely where the classical smoothness theory becomes subtle [11,22].

This distinction matters because the proofs of convergence split naturally into two parts:

- on regular patches, one can work with ordinary first differences and exploit the contraction properties of the linear scheme directly;
- near extraordinary vertices, one instead studies a finite *control ring* and the local subdivision matrix acting on that ring.

The manifold-valued proof below follows exactly the same division.

5.3 Classical concepts we will use

This section isolates the small amount of classical linear subdivision theory that the proof depends on. We do not reprove these facts here. The goal is only to set notation and to explain the role of the imported hypotheses.

5.3.1 Stationary affine-invariant schemes and local rules

Let K be a finite closed triangular or quadrilateral control mesh. A stationary linear subdivision scheme assigns to each refinement step a new control mesh by rules (called **masks**) of the form

$$\underbrace{p'_u}_{\text{refined vertex}} = \sum_{v \in \text{stencil}(u)} \alpha_{uv} \underbrace{p_v}_{\text{coarse data}}, \quad \sum_v \alpha_{uv} = 1.$$


The weights are fixed by the local combinatorial type of the stencil, not by the refinement level. Affine invariance is exactly the statement that the coefficients sum to one, so that translating or otherwise affinely transforming the control mesh commutes with refinement [11,23]. Finite support means that each refined value depends on only finitely many nearby coarse values.

The masks for the Catmull–Clark scheme are displayed in Figure 5.4 to provide a more visual example of this formula. Newly inserted vertices are called “odd” while existing vertices are called “even.” Odd vertices are further divided into “face” and “edge” vertices depending on where they are inserted. Different masks are used to handle extraordinary and boundary vertices. Masks for the Loop subdivision scheme and more details on subdivision theory can be found in [11].

The corresponding nonlinear scheme in this chapter will be obtained by keeping the same positive weights and replacing the Euclidean affine average by the Karcher mean with those weights.

5.3.2 Control rings, local subdivision matrices, and quotient spaces

Near a regular patch, contraction is naturally expressed in terms of first differences. Near an extraordinary vertex, there is no shift-invariant grid direction in which to take a single derived difference scheme. The standard replacement in the linear theory is to restrict attention to a finite *control ring* around the extraordinary vertex and to study the corresponding local subdivision ma-



Interior



Boundary

Figure 5.4: Masks for the Catmull–Clark scheme.

trix; this is the point of view used in Reif's analysis and in Zorin's exposition of subdivision near extraordinary vertices [11, 22, 24].

Fix an extraordinary valence n . Let C_n denote a finite set of vertices around the extraordinary vertex large enough that the next control ring is determined entirely by the data on this ring. In local Euclidean coordinates, let

$$E_n = \{u : C_n \rightarrow \mathbb{R}^d\}$$

be the vector space of all assignments of coordinate values to the ring. The precise dimension d is not important here; in the manifold-valued case it will be the dimension of the local normal-coordinate chart. Let

$$\mathcal{C}_n = \{u \in E_n : u(c) = a \text{ for all } c \in C_n \text{ and for some fixed } a \in \mathbb{R}^d\}$$

be the subspace of constant configurations.

The quotient space

$$E_n / \mathcal{C}_n$$

removes the constant mode of the control ring. Equivalently, two ring configurations are identified if they differ only by translating every value in the ring by the same vector. A quotient norm can be written as

$$\|[u]\|_n = \inf_{c \in \mathcal{C}_n} \|u - c\|_n,$$

so $\|[u]\|_n$ measures the distance from the ring data to the nearest constant configuration. In this sense it measures the oscillation of the control ring: constants have zero oscillation, and only the nonconstant part remains. This is the surface analogue of first differences for curves, which also vanish on constant data and ignore uniform translation.

If $B_n : E_n \rightarrow E_n$ is the local subdivision matrix on the control ring, affine invariance implies that constant configurations refine to constant configurations. Hence B_n preserves $\mathcal{C}_n$ and induces




Figure 5.5: (Blue) a particular vertex in a patch of a quadrilateral mesh. (Red) Its corresponding control ring under Catmull–Clark subdivision.

a quotient map

$$A_n : E_n/\mathcal{C}_n \rightarrow E_n/\mathcal{C}_n$$

through the relation

$$A_n \circ \pi_n = \pi_n \circ B_n,$$

where $\pi_n : E_n \rightarrow E_n/\mathcal{C}_n$ is the quotient projection. The spectral properties of A_n will supply the extraordinary-point contraction hypothesis used later.

5.3.3 Characteristic maps and imported smoothness facts

The classical linear theory of subdivision surfaces near extraordinary vertices is governed by the eigenstructure of the local subdivision matrix and by the associated characteristic map criterion of Reif [24]. Zorin’s thesis and course notes provide a clear practical exposition of this viewpoint for Loop and Catmull–Clark, emphasizing regular patches, control rings, eigenstructure, and characteristic maps [11, 22]. The nonlinear proximity-based extensions due to Wallner, Dyn, and Weinmann explain how linear regularity statements transfer to nonlinear manifold-valued analogues once a suitable proximity estimate is established [12–14].

We therefore use the following linear facts without proof:

1. On regular regions, the underlying linear Loop or Catmull–Clark scheme contracts first differences and generates the usual smooth limit surface.
2. Near each extraordinary valence occurring in the mesh, the corresponding local quotient operator A_n has spectral radius strictly less than 1.
3. Whenever the classical characteristic-map hypotheses hold for the linear scheme, the nonlinear manifold-valued analogue inherits the corresponding C^1 behavior by proximity, provided it is well defined and sufficiently close to the linear rule.

Only the first two items are needed for the continuous convergence theorem proved below. The third explains how stronger regularity results fit into the same framework, but we do not pursue them in full detail here.

5.4 From affine averages to Karcher means

Let M be a smooth Riemannian manifold with distance d . In the applications of this thesis, M is often S^3 equipped with its standard bi-invariant metric, but the present discussion does not require that specialization. The only geometric operation needed is the Karcher mean of finitely many nearby points with positive weights.

If $x_1, \dots, x_m \in M$ lie inside a sufficiently small strongly convex geodesic ball, and if the weights $\alpha_1, \dots, \alpha_m \geq 0$ satisfy $\sum_i \alpha_i = 1$, then their weighted Karcher mean is the unique minimizer of

$$x \mapsto \frac{1}{2} \sum_{i=1}^m \alpha_i d(x, x_i)^2.$$

A convenient sufficient condition for existence and uniqueness is that all stencil points lie in a ball whose radius is below the standard curvature–injectivity threshold; see Afsari for a precise statement and for the convex-hull property of the resulting mean [25]. This is the small-data hypothesis that will appear again in the proof.

The manifold-valued subdivision rule T is obtained by replacing each positive affine average in the linear scheme S with the corresponding Karcher mean. Thus, whenever the linear rule reads

$$\underbrace{p'_u}_{\text{Euclidean child value}} = \sum_v \alpha_{uv} p_v,$$

the manifold-valued rule becomes

$$\underbrace{P'_u}_{\text{manifold-valued child value}} = \text{KM}(\{(\alpha_{uv}, P_v) : v \in \text{stencil}(u)\}),$$

where KM denotes the weighted Karcher mean of the stencil values P_v with weights α_{uv} .

The replacement is local, coordinate-free, and agrees with the Euclidean rule to first order when the stencil diameter is small. That first-order agreement is the source of the proximity estimate used in the proof.

5.5 Proof roadmap

The proof is easiest to digest if one keeps in mind exactly what is imported from classical theory and what is proved here.

Classical subdivision facts used without proof

- regular-region contraction for the chosen linear subdivision scheme;
- spectral contraction of the local quotient operator near each extraordinary valence;
- the broader classical characteristic-map framework, cited only for context.

New arguments supplied in this chapter

- the standard local quadratic expansion of the Karcher mean near constant data, recorded in the form needed for the proof;
- a one-step proximity estimate between the nonlinear rule T and the linear rule S ;
- perturbative contraction on regular patches;
- perturbative contraction of extraordinary control-ring oscillations modulo constants;
- geometric decay of a global mesh-size functional;
- uniform convergence of continuous interpolants to a continuous limit surface.

This split matters because the proof relies on standard contraction and characteristic-map facts from the linear subdivision literature; the new part is the proximity argument showing that the Karcher-mean rule remains close enough to the linear rule in the small-data regime. This is almost the same logic as the curve proof of Wallner and Dyn [12]. The one substantial change is that, for surfaces, first differences on a shift-invariant grid are no longer the right object near extraordinary vertices. They are replaced by oscillations of a finite control ring modulo constants, encoded by the quotient space $E_n/\mathcal{C}_n$.

5.6 Continuous convergence by proximity

We now state the theorem proved in this chapter. The statement is intentionally abstract: it applies to any stationary affine-invariant linear subdivision surface scheme with positive masks

once the listed hypotheses are verified. Loop and Catmull–Clark are the concrete schemes to keep in mind.

5.6.1 Hypotheses

Let M be the Riemannian manifold in which the data take values. Let S be a stationary affine-invariant linear subdivision surface scheme on a fixed finite triangular or quadrilateral mesh K , with finite support and nonnegative masks. Let T be the manifold-valued analogue obtained by replacing positive affine averages with Karcher means carrying the same weights. We write $|K|$ for the underlying topological surface of the initial mesh.

We assume the following.

(H1) **Strongly convex stencil neighborhoods.** There exists $\rho > 0$ such that every stencil arising during refinement lies inside a strongly convex geodesic ball of radius $< \rho$. Hence all Karcher means used by the scheme are uniquely defined.

(H2) **Regular-patch contraction for the linear scheme.** There exists $0 < \mu_{\text{reg}} < 1$ such that

$$\delta_{\text{reg}}(S(P)) \leq \mu_{\text{reg}} \delta_{\text{reg}}(P)$$

on each regular patch, where $\delta_{\text{reg}}(P)$ denotes the maximal regular first difference.

(H3) **Extraordinary-ring contraction for the linear scheme.** For each extraordinary valence n occurring in K , there exists a norm $\|\cdot\|_n$ on $E_n/\mathcal{C}_n$ and a constant $0 < \mu_n < 1$ such that

$$\|A_n \xi\|_n \leq \mu_n \|\xi\|_n \quad \text{for all } \xi \in E_n/\mathcal{C}_n.$$

(H4) **Small-data regime.** The initial control mesh P^0 is dense enough that all stencils remain in the strongly convex neighborhoods of (H1) and in normal-coordinate neighborhoods where the Taylor estimates used in Lemma 5.6.1 and Corollary 5.6.1 are uniform.

The next theorem is the formal statement of Theorem 1.4.2 from the introduction.

Theorem 5.6.1 (Continuous convergence of manifold-valued subdivision).

Assume (H1)–(H4). Then the sequence of manifold-valued refined control meshes

$$P^j := T^j P^0$$

admits continuous interpolants $I_j : |K| \rightarrow M$ which converge uniformly to a continuous limit map

$$I_\infty : |K| \rightarrow M.$$

Equivalently, the manifold-valued subdivision rule T generates a continuous limit surface.

The proof is organized into six short steps.

5.6.2 Local expansion of the Karcher mean

The key nonlinear fact is that the Karcher mean agrees to first order with the corresponding Euclidean affine average. This is a standard consequence of the smooth dependence of the Riemannian center of mass on nearby data; see Karcher and Afsari for the general theory of such means [7, 25]. We reproduce here the estimate in the particular form needed for the proximity proof.

Lemma 5.6.1 (Quadratic expansion near constant data).

Fix nonnegative weights $\alpha_1, \dots, \alpha_m$ with $\sum_i \alpha_i = 1$. Let U be a compact subset of a strongly convex normal-coordinate chart. Then there exist constants $r > 0$ and $C > 0$ such that the following holds.

For any $c \in U$ and any tangent vectors $u_1, \dots, u_m \in T_c M$ with $\|u_i\| \leq r$, let

$$x_i = \exp_c(u_i).$$

Then the Karcher mean $G_\alpha(x_1, \dots, x_m)$ can be written as

$$G_\alpha(x_1, \dots, x_m) = \exp_c \left(\underbrace{\sum_{i=1}^m \alpha_i u_i}_{\text{Euclidean affine average}} + \underbrace{R_c(u_1, \dots, u_m)}_{\text{quadratic remainder}} \right),$$

with

$$\|R_c(u_1, \dots, u_m)\| \leq C \left(\max_i \|u_i\| \right)^2.$$

Proof. Fix $c \in U$ and consider the smooth map

$$F_c(z, u_1, \dots, u_m) = \sum_{i=1}^m \alpha_i \log_{\exp_c(z)}(\exp_c(u_i))$$

for z and the u_i sufficiently small. The Karcher mean is characterized by the equation

$$F_c(z, u_1, \dots, u_m) = 0.$$

At the diagonal configuration $u_1 = \dots = u_m = 0$, the solution is $z = 0$. Moreover,

$$D_z F_c(0, 0, \dots, 0) = -I, \quad D_{u_i} F_c(0, 0, \dots, 0) = \alpha_i I.$$

The implicit function theorem therefore gives a unique smooth solution

$$z = G_c(u_1, \dots, u_m)$$

near the origin. Differentiating the identity

$$F_c(G_c(u), u) = 0$$

at $u = 0$ yields

$$DG_c(0)(v_1, \dots, v_m) = \sum_{i=1}^m \alpha_i v_i.$$

Thus the first-order Taylor polynomial of G_c is the Euclidean affine average, and the remainder is quadratic. Since U is compact and the chart geometry varies smoothly with c , the constants can be chosen uniformly in c . This gives the claimed estimate. $\square$

Corollary 5.6.1 (Finite compositions of intrinsic averages).

Any finite composition of Karcher averaging maps with fixed nonnegative weights agrees, in local coordinates and to first order at constant data, with the corresponding affine averaging composition. Consequently, on each stencil type, the nonlinear rule T is a quadratic perturbation of the linear rule S .

Proof. Apply Lemma 5.6.1 to each averaging step. Because only finitely many steps occur in each stencil, the sum of the Taylor remainders remains $O(\sigma^2)$ when the input stencil diameter is $O(\sigma)$. $\square$

5.6.3 One-step proximity

To encode all local oscillations in one quantity, define

$$\sigma(P) = \max \left\{ \underbrace{\delta_{\text{reg}}(P)}_{\text{regular first differences}}, \underbrace{\max_{v \in \text{Ext}(K)} \|\eta_v(P)\|_{n(v)}}_{\text{extraordinary-ring oscillations modulo constants}} \right\}.$$

Here $\eta_v(P) \in E_{n(v)}/\mathcal{C}_{n(v)}$ is the quotient class of the chosen control ring near the extraordinary vertex v .

Lemma 5.6.2 (Global one-step proximity).

There exists $C > 0$ such that for every sufficiently dense control mesh P ,

$$\|T(P) - S(P)\|_{\infty} \leq C \sigma(P)^2$$

in local coordinates, stencil by stencil.

Proof. There are only finitely many stencil types on the fixed finite mesh topology: finitely many regular stencils and finitely many extraordinary valences. On each type, Corollary 5.6.1 gives

$$T_{\text{stencil}} = S_{\text{stencil}} + O(\sigma^2).$$

Taking the maximum over this finite collection yields the claim. $\square$

Lemma 5.6.3 (Displacement safety).

There exists $C > 0$ such that every child vertex w produced from a control mesh P lies within $C \sigma(P)$ of the vertices of its parent face.

Proof. For the linear scheme S , every child vertex is a convex combination of finitely many nearby parent vertices because the masks are nonnegative and sum to one. Hence the linear child lies in the Euclidean convex hull of its parent stencil, so its distance from any parent-face vertex is $O(\sigma(P))$. By Lemma 5.6.2, the nonlinear child differs from the linear child by $O(\sigma(P)^2)$, which is a lower-order perturbation once $\sigma(P)$ is sufficiently small. $\square$

5.6.4 Contraction on regular patches

Lemma 5.6.4 (Regular-patch perturbative contraction).

There exists $C_{\text{reg}} > 0$ such that

$$\delta_{\text{reg}}(T(P)) \leq \mu_{\text{reg}} \delta_{\text{reg}}(P) + C_{\text{reg}} \sigma(P)^2.$$

Proof. Take a regular refined edge with endpoints u and v . Then

$$d(T(P)_u, T(P)_v) \leq d(S(P)_u, S(P)_v) + d(T(P)_u, S(P)_u) + d(T(P)_v, S(P)_v).$$

Taking the maximum over all regular refined edges and applying Lemma 5.6.2 gives

$$\delta_{\text{reg}}(T(P)) \leq \delta_{\text{reg}}(S(P)) + 2C \sigma(P)^2 \leq \mu_{\text{reg}} \delta_{\text{reg}}(P) + C_{\text{reg}} \sigma(P)^2.$$

□

5.6.5 Extraordinary-ring contraction

Fix an extraordinary valence n . Let

$$\Phi_n : E_n \rightarrow E_n$$

be the local nonlinear refinement map written in local coordinates around constant data. Since constant data refine to constant data, Φ_n preserves $\mathcal{C}_n$ and therefore induces a quotient map

$$\widehat{\Phi}_n : E_n/\mathcal{C}_n \rightarrow E_n/\mathcal{C}_n.$$

Lemma 5.6.5 (Linearization at constant data).

The derivative of the quotient map at the origin equals the linear quotient operator:

$$D\widehat{\Phi}_n(0) = A_n.$$

Proof. At constant data, each Karcher averaging step has the same first derivative as the corresponding affine averaging step by Lemma 5.6.1. Therefore the derivative of the full nonlinear local rule equals the linear subdivision matrix B_n . Passing to the quotient by constants gives the claim. □

Lemma 5.6.6 (Quadratic remainder on the quotient).

There exists $C_n > 0$ such that

$$\|\widehat{\Phi}_n(\xi) - A_n \xi\|_n \leq C_n \|\xi\|_n^2$$

for all ξ sufficiently small.

Proof. This follows from the Taylor remainder estimate at the origin for $\widehat{\Phi}_n$, using Lemma 5.6.5. □

Proposition 5.6.1 (Extraordinary-ring contraction).

For each extraordinary valence n , there exist $q_n \in (0, 1)$ and $\varepsilon_n > 0$ such that

$$\|\widehat{\Phi}_n(\xi)\|_n \leq q_n \|\xi\|_n \quad \text{whenever } \|\xi\|_n \leq \varepsilon_n.$$

Proof. By (H3),

$$\|A_n \xi\|_n \leq \mu_n \|\xi\|_n \quad (\mu_n < 1).$$

By Lemma 5.6.6,

$$\|\widehat{\Phi}_n(\xi)\|_n \leq \mu_n \|\xi\|_n + C_n \|\xi\|_n^2.$$

Choose $\varepsilon_n > 0$ so small that

$$q_n := \mu_n + C_n \varepsilon_n < 1.$$

Then the stated inequality follows. □

5.6.6 Global decay and convergence of interpolants

Proposition 5.6.2 (Global contraction of the size functional).

There exist $\varepsilon > 0$ and $q \in (0, 1)$ such that whenever $\sigma(P) \leq \varepsilon$,

$$\sigma(T(P)) \leq q \sigma(P).$$

Proof. For the regular part, Lemma 5.6.4 gives

$$\delta_{\text{reg}}(T(P)) \leq \mu_{\text{reg}} \delta_{\text{reg}}(P) + C_{\text{reg}} \sigma(P)^2 \leq (\mu_{\text{reg}} + C_{\text{reg}}\varepsilon)\sigma(P).$$

For the extraordinary part, Proposition 5.6.1 yields

$$\|\eta_v(T(P))\|_{n(v)} \leq q_{n(v)} \|\eta_v(P)\|_{n(v)}$$

provided $\sigma(P)$ is below all ε_n . Since only finitely many extraordinary valences occur in the fixed mesh K , choose $\varepsilon > 0$ so small that

$$q = \max \left\{ \mu_{\text{reg}} + C_{\text{reg}}\varepsilon, \max_n q_n \right\} < 1.$$

Then each term entering $\sigma(T(P))$ is bounded by $q\sigma(P)$. □

Corollary 5.6.2.

If $\sigma(P^0) \leq \varepsilon$, then

$$\sigma(P^j) \leq q^j \sigma(P^0) \quad \text{for all } j \geq 0.$$

Proof. Immediate by induction from Proposition 5.6.2. □

To pass from control meshes to an actual surface, attach to each refined mesh a continuous interpolant. For triangular faces, use barycentric Karcher interpolation of the three vertex values. For quadrilateral schemes, first fix an auxiliary triangulation of each face; this triangulation is used only to define the interpolant and does not alter the refinement rule itself. Because shared

edges correspond to the same one-dimensional geodesic interpolation, the facewise construction is continuous across edges.

Lemma 5.6.7 (Uniform step-to-step difference).

There exists $C > 0$ such that the interpolants $I_j : |K| \rightarrow M$ satisfy

$$\|I_{j+1} - I_j\|_\infty \leq C \sigma(P^j) \quad \text{for all } j.$$

Proof. Fix a parent face F at level j and a child triangle $\tau \subset F$ at level $j + 1$. By Corollary 5.6.2, the vertices of F are pairwise $O(\sigma(P^j))$ -close. By Lemma 5.6.3, every child vertex inside F is also $O(\sigma(P^j))$ -close to the parent-face vertices. Hence all parent and child vertices associated with τ lie in a common strongly convex ball of diameter $O(\sigma(P^j))$.

The image of a geodesic barycentric interpolation remains inside the strongly convex hull of its vertices; see [25]. Therefore both $I_j(\tau)$ and $I_{j+1}(\tau)$ stay inside a set of diameter $O(\sigma(P^j))$, giving

$$\sup_{x \in \tau} d(I_{j+1}(x), I_j(x)) \leq C \sigma(P^j).$$

Since only finitely many face types occur, the constant may be chosen uniformly. □

Proof of Theorem 5.6.1. By Lemma 5.6.7 and Corollary 5.6.2,

$$\|I_{j+1} - I_j\|_\infty \leq C \sigma(P^j) \leq C q^j \sigma(P^0).$$

Because $0 < q < 1$, the series

$$\sum_{j=0}^{\infty} C q^j \sigma(P^0)$$

converges. Hence (I_j) is uniformly Cauchy and therefore converges uniformly to a limit map $I_\infty : |K| \rightarrow M$. Each I_j is continuous, so the uniform limit I_∞ is continuous as well. □

5.7 Consequences and role in the thesis

The theorem proved here has a limitation: it guarantees continuous convergence, but smoothness is not proved. Smoothness is a separate question even for Euclidean subdivision schemes, where it depends on additional spectral and characteristic-map hypotheses. The manifold-valued construction should therefore not be expected to be smooth merely from the convergence proof above. Proving C^1 or higher regularity would require importing the corresponding linear smoothness theory and proving a stronger proximity estimate.

Still, it suggests that the manifold-valued subdivision rule T is a viable multiresolution refinement method for the auxiliary data in the retargeting pipeline. The coarse-to-fine optimization can be performed by first solving on a coarse mesh and then refining both positions and manifold-valued data together, with the guarantee that the refined data remain close to the linear template and therefore continue to interact well with the energy design. The practical consequence is modest but useful: if a coarse solution is refined, the quaternion field can be refined by the same subdivision stencil as the positions, replacing affine averages by Karcher means while preserving continuous convergence in the small-data regime.

Chapter 6

Illustrative Applications and Experimental Formulations

This chapter describes three illustrative case studies for the framework developed in the preceding chapters. The emphasis is partly on formulation and partly on numerical behavior: each section specifies the target geometry, the optimization variables, the energy to be minimized, and the observed convergence of the implementation used to generate the accompanying figures.

The case studies are deliberately minimal but complementary. The first fits a shirt around an already aligned body and isolates the role of one-sided clearance together with near-isometry. The second adds semantic landmark information to a garment under pose change. The third changes the search space itself: the mesh is no longer deformed freely, but only through an armature or curve-driven control model. These examples are organized around three questions:

1. Can a near-isometric energy resolve clipping while preserving designed geometric structure?
2. Can landmark constraints and regionwise compliance be added without changing the essential method?
3. Can the same mesh-space energy be pulled back through a lower-dimensional control model when direct vertex optimization is undesirable?

6.1 Common residuals and solver conventions

6.1.1 A smooth clearance residual

Let $d_i \geq 0$ denote a desired clearance or offset at vertex i , and define the smooth hinge

$$\psi_\beta(t) = \frac{1}{\beta} \log(1 + e^{\beta t}), \quad \beta > 0.$$

Equivalently, in the implementation we write the smoothing width as $\varepsilon = \beta^{-1}$ and use

$$\rho_\varepsilon(t) = \varepsilon \log(1 + e^{t/\varepsilon}).$$

When $t \ll 0$, this quantity is close to 0; when $t \gg 0$, it is close to t . The associated one-sided clearance term is

$$\mathcal{E}_{\text{clear}}(\mathbf{P}) = \sum_{i \in \mathcal{V}} \rho_\varepsilon(d_i - \widehat{\Phi}_{\text{SDF}}(\mathbf{p}_i))^2 \text{area}(\star i).$$

This is the smooth replacement for a hard inequality constraint. It is not the only possible choice, but it keeps the objective compatible with the smooth Gauss–Newton and Levenberg–Marquardt viewpoint discussed earlier.

When attraction to a prescribed offset level set is desired, we also use

$$\mathcal{E}_{\text{off}}(\mathbf{P}) = \sum_{i \in \mathcal{V}} (\widehat{\Phi}_{\text{SDF}}(\mathbf{p}_i) - d_i)^2 \text{area}(\star i).$$

The first term is one-sided and enforces exteriority; the second is two-sided and attracts the mesh toward the desired level set. In the loaded shirt examples below, the offset attraction is disabled, because loose regions of the shirt should not be pulled onto the body merely because they are farther than the desired clearance.

6.1.2 A weighted near-isometric term

To allow for spatially varying rigidity, the ARAP-like term of Chapter 2 is used in the weighted form

$$\mathcal{E}_{\text{ARAP}}^\kappa(\mathbf{P}, \mathbf{Q}) = \sum_{i \in \mathcal{V}} \kappa_i \sum_{j \in \mathcal{N}_1(i)} \mathbf{w}_{ij} \left\| q_i(\mathbf{p}_i^{(0)} - \mathbf{p}_j^{(0)}) q_i^* - (\mathbf{p}_i - \mathbf{p}_j) \right\|^2 \text{area}(\star i),$$

where $\kappa_i \geq 0$ is a local rigidity weight. Large values of κ_i enforce a stronger near-isometric deformation, while smaller values allow more compliance in selected regions.

6.1.3 Experimental protocol and reported quantities

For each case study we report the size of the conforming and target meshes, the target-field construction, the active energy weights, and the numerical solver settings. The objectives are nonconvex, so convergence should be read as convergence to a local stationary point or to a stable feasible configuration, not as a certificate of global optimality.

The clearance diagnostics use two simple quantities. The maximum clearance violation is

$$\max_i \max(0, d_i - \Phi(\mathbf{p}_i)),$$

where Φ denotes the target signed-distance query used in that experiment. The violating-vertex count is the number of vertices for which this one-sided violation is positive above the stopping tolerance. Thus the maximum violation measures the depth of the worst remaining interpenetration, while the vertex count measures how much of the mesh is still affected.

The plots shown below normalize energy, maximum violation, and violating-vertex count so that their trends can be compared in one diagram. Their vertical scale is therefore not a physical unit scale. The tables give the corresponding unnormalized values. Iteration was terminated after the mesh was feasible for a trailing window of samples and the energy decrease over a 16-sample window fell below the reported tolerance. Unless otherwise stated, the feasibility tolerance was 10^{-4} for the maximum violation and zero violating vertices.

All experiments were run on an AMD Ryzen 7 2700X CPU and an NVIDIA GeForce RTX 2080 Ti GPU. ADSDF point queries, as well as the preview renderer used to generate the figures, run through CUDA. The current implementation uses single precision for geometry and field queries and double precision for accumulated scalar diagnostics and the small dense or sparse linear algebra routines used in LM refinements.

6.2 Shirt fitting on an aligned body

6.2.1 Geometric setup and intended phenomenon

The first example is a shirt fitting problem. The target is a body mesh and the conforming mesh is a shirt already placed in roughly the correct coordinate frame. The initial placement is close enough that a rigid alignment problem is not the main issue, but several shirt vertices penetrate the body. The desired behavior is modest: resolve the clipping while changing the shirt as little as possible.

This case isolates the contact-preserving role of the clearance term. It also demonstrates why the distance term is not always appropriate for garments. A two-sided offset term would pull loose shirt regions toward the body everywhere. For this example, only the one-sided clearance term, the ARAP term, and a weak center stabilizer are used.

6.2.2 Optimization variables and energies

Let $\widehat{\Phi}_{\text{body}}$ denote the ADSDF of the body. It is built from the body mesh on a degree-2 grid using a linear coefficient filter, with ray-parity mesh signing used during construction. The optimization variables are the shirt vertices

$$\mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_N) \in (\mathbb{R}^3)^N$$

and the auxiliary quaternion field

$$\mathbf{Q} = (q_1, \dots, q_N) \in (S^3)^N.$$

The reference mesh is the initial shirt placement $\mathbf{P}^{(0)} = \mathbf{P}_{\text{init}}$, and the desired clearance is constant:

$$d_i \equiv 0.025.$$

Quantity	Value
conforming mesh	266 vertices, 483 triangles after triangulation
target mesh	7534 vertices, 15064 triangles after triangulation
target field	ADSDF built from the body mesh
ADSDF grid	$56 \times 40 \times 64$
ADSDF polynomial degree	2
ADSDF filter	linear coefficient filter
sign convention	ray-parity mesh signing
desired clearance d	0.025
softplus width ε	0.008
active weights	$T_{\text{clear}} = 120, T_{\text{ARAP}} = 1, T_{\text{ctr}} = 0.02$
disabled weights	$T_{\text{off}} = 0$
solver	gradient descent with backtracking line search
nominal step sizes	position 0.012, rotation 0.02
maximum iterations	220
stopping criterion	feasible trailing window and 16-sample energy decrease $\leq 10^{-3}$
accepted samples	21
wall-clock runtime	1047.8 ms

Table 6.1: Parameters for the aligned body/shirt clearance experiment.

Metric	Initial	Final
total energy	0.193751	0.00130927
maximum clearance violation	0.103320	0
violating vertices	156	0
minimum signed distance	-0.0783201	0.0319435
mean time between samples after ADSDF setup	1.47 ms	

Table 6.2: Numerical behavior of the aligned body/shirt clearance experiment.

The energy is

$$\mathcal{E}_{\text{fit}}(\mathbf{P}, \mathbf{Q}) = T_{\text{clear}}\mathcal{E}_{\text{clear}}(\mathbf{P}) + T_{\text{ARAP}}\mathcal{E}_{\text{ARAP}}^{\kappa}(\mathbf{P}, \mathbf{Q}) + T_{\text{ctr}}\mathcal{E}_{\text{ctr}}(\mathbf{P}).$$

The weights were chosen by a small manual tuning pass with a concrete criterion: the final shirt should have no violating vertices and should not visibly collapse toward the body. This led to a relatively strong clearance weight, a unit ARAP weight, and only a weak center term.

Figure 6.1 shows the qualitative effect of the optimization. The numerical values in Table 6.2 show the same behavior quantitatively: all 156 initially violating vertices are cleared, and the




Figure 6.1: A shirt fitting example on the loaded body and shirt meshes. The initial shirt penetrates the body. The optimized result clears the body while preserving the overall shirt shape.

worst violation drops to zero. After ADSDF setup, the recorded per-sample cost is only a few milliseconds for this mesh size.

6.3 Pose-retargeted shirt with landmarked openings

6.3.1 Geometric setup and intended phenomenon

The second example keeps the source asset soft rather than hard and makes the role of landmarks explicit. A shirt mesh is given on one body and must be adapted to a target body in a different pose. The desired deformation is not captured by contact alone: the neck opening, lower hem, and sleeve openings should move to semantically appropriate places on the target body while the shirt panels remain as faithful as possible to the original garment.

The intended phenomenon is therefore different from the previous case. The example is not about showing that landmarks are unnecessary, but precisely the opposite: certain semantic features of the source garment are best controlled by designated target features. At the same time, the shirt should remain relatively stiff in most regions and only become more compliant where the pose change demands it, for example near the shoulders.

6.3.2 Landmark formulation

Let

$$\mathcal{L} = \{\text{neck, hem, left sleeve, right sleeve}\}$$




Figure 6.2: Diagnostics for the aligned body/shirt clearance experiment. The curves show normalized total energy, maximum clearance violation, and violating-vertex count over the active convergence window.

denote the four boundary loops we intend to control. For each $\ell \in \mathcal{L}$, let $B_\ell \subset \mathcal{V}$ be the corresponding set of boundary vertices on the shirt mesh, and let Γ_ℓ be the target curve or target region prescribed on the posed body. In the implemented example, the boundary vertices of the original shirt are used as landmark vertices, and the target positions are read by matching vertex index from a manually edited posed shirt. The interior of that posed shirt is ignored; it is used only as a convenient way to specify target boundary locations.

With point correspondences along the loops, the landmark term is

$$\mathcal{E}_{\text{loop}}(\mathbf{P}) = \sum_{\ell \in \mathcal{L}} \sum_{i \in B_\ell} \omega_i^\ell \|\mathbf{p}_i - \mathbf{p}_i^\star\|^2.$$

Before optimization, a harmonic displacement seed is computed on the shirt graph. The boundary displacements are fixed to the landmark targets, and the interior displacements are smoothed over 480 iterations. This is a simple programmatic analogue of weight painting: the boundary features are placed where the artist requested, while the interior receives a smooth initial displacement rather than an abrupt local edit.




Figure 6.3: Landmark-guided shirt retargeting. (Left) A red shirt asset fitted to a T-posed body, with landmark features visualized in blue. (Right) The body pose and landmark positions are modified before optimization.

6.3.3 Optimization problem

Let $\hat{\Phi}_{\text{pose}}$ denote the ADSDF of the posed target body. We again optimize over vertex positions and quaternion data:

$$(\mathbf{P}, \mathbf{Q}) \in (\mathbb{R}^3)^N \times (S^3)^N.$$

The total energy is

$$\mathcal{E}_{\text{shirt}}(\mathbf{P}, \mathbf{Q}) = T_{\text{clear}}\mathcal{E}_{\text{clear}}(\mathbf{P}) + T_{\text{ARAP}}\mathcal{E}_{\text{ARAP}}^{\kappa}(\mathbf{P}, \mathbf{Q}) + T_{\text{land}}\mathcal{E}_{\text{loop}}(\mathbf{P}) + T_{\text{ctr}}\mathcal{E}_{\text{ctr}}(\mathbf{P}).$$

The offset term is again disabled. The ARAP weight is handled by continuation: it starts at 80 to prevent early collapse of the sleeves and then relaxes by factors of 0.5 when progress stalls, down to a final value of 2. This is a practical tuning heuristic rather than an optimal parameter-selection rule. It reflects the intended modeling behavior: preserve the garment strongly at first, then allow more compliance only after the landmarks and clearance term have made the main pose correction.

Figure 6.3 emphasizes the role of the landmark information. Without such a term, clearance and ARAP can remove penetration but do not determine where semantic boundary features such

Quantity	Value
conforming mesh	266 vertices, 483 triangles after triangulation
target mesh	7534 vertices, 15064 triangles after triangulation
landmark vertices	59 boundary vertices
ADSDF grid	$56 \times 40 \times 64$
ADSDF polynomial degree	2
ADSDF filter	linear coefficient filter
LSQ stencil	fine radius 2, fine extent 0.45, regularization 10^{-7}
sign convention	ray-parity mesh signing
desired clearance d	0.025
softplus width ε	0.008
clearance weight	5000
landmark weight	1.2
center weight	0.02
ARAP continuation	80, 40, 20, 10, 5, 2.5, 2
disabled weights	$T_{\text{off}} = 0$
solver	gradient descent with backtracking line search
nominal step sizes	position 0.004, rotation 0.02
maximum iterations	720
stopping criterion	final ARAP stage, feasible trailing window, and 16-sample energy decrease $\leq 9.50392 \cdot 10^{-5}$
accepted samples	276
wall-clock runtime	2294.4 ms

Table 6.3: Parameters for the ADSDF body/shirt landmark retargeting experiment.

as cuffs and the neckline should move. In the reported run, the harmonic seed first transfers the boundary landmarks smoothly into the shirt interior, and the subsequent optimization resolves clearance while allowing the landmark error to increase only slightly. Table 6.4 shows that the final result has no clearance violations and a boundary landmark RMS error of about 0.007 in the mesh coordinate system.

6.4 Armature-driven tentacle in a synthetic obstacle field

6.4.1 Geometric setup and intended phenomenon

The final example changes the search space rather than just the energy. A tentacle-like mesh is guided by a low-dimensional armature and placed near a synthetic static obstacle field. The initial pose is intended to be roughly correct, but local clipping remains. The goal is not to alter

Metric	Seeded initial	Final
total energy	9.50392	0.00884370
maximum clearance violation	0.141042	0
violating vertices	96	0
minimum ADSDF value	-0.116042	0.0286324
boundary landmark RMS	$1.23 \cdot 10^{-8}$	0.00701227
mean time between samples after ADSDF setup	1.54 ms	

Table 6.4: Numerical behavior of the ADSDF body/shirt landmark retargeting experiment. The initial row is measured after the harmonic displacement seed. The raw unseeded shirt has boundary landmark RMS error about 0.450 before retargeting.




Figure 6.4: Diagnostics for the ADSDF body/shirt landmark retargeting experiment. The plotted curves show the seeded optimization after the harmonic landmark initialization.

the mesh arbitrarily; instead, the optimizer should change the control parameters just enough to resolve contact, while preserving the resulting mesh geometry as much as possible.

This is intended to demonstrate two things at once. First, the framework can interact with geometrically nontrivial static targets. Second, the energy can be pulled back through a control model, so that the method corrects the pose mechanism without abandoning the mesh-based geometric prior. Self-collision of the tentacle is not treated in this experiment.

6.4.2 Armature-restricted search space

Let

$$\theta \in \Theta = G_1 \times \cdots \times G_B$$

denote the armature parameters, where each factor G_b is the configuration space of one control element. The precise parameterization is not important here; what matters is that the pose induces a deformed mesh

$$\mathbf{P}(\theta) = (\mathbf{p}_1(\theta), \dots, \mathbf{p}_N(\theta)).$$

For a standard weighted-bone model, one may write

$$\mathbf{p}_i(\theta) = \sum_{b=1}^B \omega_{ib} T_b(\theta) \bar{\mathbf{p}}_i,$$

where $\bar{\mathbf{p}}_i$ is the rest-pose vertex, $T_b(\theta)$ is the transform associated with bone b , and ω_{ib} are the skinning weights. The implemented prototype uses 9 centerline control points, with endpoints pinned, and programmatically generated smooth compact weights for the tube vertices.

6.4.3 Energy and numerical problem

Let $\hat{\Phi}_{\text{obj}}$ denote the ADSDF of the static obstacle field. We retain the auxiliary quaternion field

$$\mathbf{Q} = (q_1, \dots, q_N) \in (S^3)^N$$

because the near-isometric prior is still expressed in terms of local mesh rotations rather than armature variables.

A simple mesh-space guidance term is imposed on the control curve. If θ^{guide} denotes the artist-prescribed guide controls, define

$$\mathcal{E}_{\text{guide}}(\theta) = \sum_b \left\| \theta_b - \theta_b^{\text{guide}} \right\|^2.$$

Quantity	Value
conforming mesh	2628 vertices, 5184 triangles
target obstacle mesh	1952 vertices, 3888 triangles
control variables	9 centerline controls, endpoints pinned
ADSDF grid	$48 \times 40 \times 40$
ADSDF polynomial degree	2
ADSDF filter	linear coefficient filter
LSQ stencil	fine radius 2, fine extent 0.55, regularization 10^{-7}
sign convention	ray-parity mesh signing
desired clearance d	0.055
softplus width ε	0.020
active weights	$T_{\text{clear}} = 360, T_{\text{ARAP}} = 0.08, T_{\text{guide}} = 0.020, T_{\text{bend}} = 0.04$
solver	gradient descent with control-space LM refinement
initial LM damping	10^{-3} , decreased by 0.35 on accepted LM steps and increased by 8 on rejection
maximum gradient iterations	300
stopping criterion	feasible trailing window and 16-sample energy decrease ≤ 0.0100123
accepted samples	122
wall-clock runtime	2563.6 ms

Table 6.5: Parameters for the ADSDF armature-driven tentacle experiment.

A separate bending term penalizes changes in the second difference of the control curve. The full objective is then

$$\mathcal{E}_{\text{tent}}(\theta, \mathbf{Q}) = T_{\text{clear}}\mathcal{E}_{\text{clear}}(\mathbf{P}(\theta)) + T_{\text{ARAP}}\mathcal{E}_{\text{ARAP}}^{\kappa}(\mathbf{P}(\theta), \mathbf{Q}) + T_{\text{guide}}\mathcal{E}_{\text{guide}}(\theta) + T_{\text{bend}}\mathcal{E}_{\text{bend}}(\theta).$$

The positional gradient is pulled back by the rig Jacobian:

$$\text{grad } \mathcal{E}_{\text{tent}}(\theta) = (DS(\theta))^* \text{grad}_{\mathbf{P}} \mathcal{E}_{\text{tent}}(\mathbf{P}(\theta), \mathbf{Q}).$$

Figure 6.5 shows the role of the restricted control space. The optimized mesh is not an arbitrary free-form deformation; it is the mesh induced by the optimized control curve. Table 6.6 shows that the largest clearance violation decreases from about 0.379 to zero, while the guide-control RMS increases to about 0.396. This is the intended tradeoff: the guide term keeps the armature near the artist’s pose, but it is not a hard constraint.

Metric	Initial	Final
total energy	10.0123	0.0199382
maximum clearance violation	0.379297	0
violating vertices	1769	0
minimum ADSDF value	-0.324297	0.189989
guide-control RMS	0	0.396017
mean time between samples	13.34 ms	

Table 6.6: Numerical behavior of the ADSDF armature-driven tentacle experiment.




Figure 6.5: Armature-driven tentacle interaction with a synthetic obstacle field. The optimization removes the initial interpenetration by moving the low-dimensional control curve, rather than by freely editing every mesh vertex.

6.5 Robustness and computational cost

The three examples above all converge to configurations with zero violating vertices under their respective signed-distance queries. This should not be read as a general robustness theorem. The examples are deliberately small, and the weights were chosen by a limited manual tuning process guided by the qualitative goal of clearing visible clipping without destroying the source shape. The results do, however, show that the same residual vocabulary can be used across direct mesh fitting, landmark-guided retargeting, and armature-restricted deformation.

The per-sample costs are small relative to ordinary interactive time scales: about 1.47 ms for aligned shirt fitting after ADSDF setup, 1.54 ms for the ADSDF shirt retargeting run after ADSDF setup, and 13.34 ms for the ADSDF tentacle run. The full wall-clock times include setup, field construction where applicable, line-search evaluations, diagnostic logging, and preview rendering. In an interactive tool, the target field would typically be built once and reused across many local




Figure 6.6: Diagnostics for the ADSDF armature-driven tentacle experiment. The maximum violation and violating-vertex count drop to zero early, after which the energy decreases as the armature settles within the remaining guide and preservation terms.

edits. On these examples, roughly real-time feedback therefore appears plausible, although the present prototype is not yet an optimized interactive application.

The main sensitivity observed in these runs is the balance between clearance and preservation. If the clearance weight is too small, intersections can remain. If the preservation weight is too strong throughout the run, pose changes such as the shirt sleeves can be resisted too much. The continuation used in the landmarked shirt example is a practical response to this issue: start with a strong ARAP weight to prevent early collapse, and relax it only after the main landmark and clearance progress stalls.

Chapter 7

Conclusion

7.1 Summary of contributions

This thesis developed a modular framework for mesh retargeting against target surfaces represented by signed distance fields. The central idea was to combine target-side geometric information from a signed distance representation with deformation energies designed to preserve design intent. In this framework, the target contributes quantities such as signed distance, normals, and, when needed, Hessian-derived curvature information; the source mesh contributes Euclidean vertex positions together with manifold-valued rotational data. The resulting optimization problem accommodates both direct mesh retargeting and constrained variants, such as rig-driven or armature-driven formulations.

Three contributions are especially central.

1. The first is the overall *SDF-plus-energies* viewpoint: represent the target by a differentiable signed distance field, then formulate retargeting as the minimization of a modular energy built from clearance, alignment, near-isometry, landmarks, and related terms. This separates target representation, energy design, and optimization in a way that makes the framework extensible. We demonstrated this extensibility in multiple scenarios: direct shirt fitting against an aligned body, landmark-driven shirt retargeting under a pose change, and armature-driven retargeting of a tentacle in a synthetic obstacle field that allowed us to operate on a lower-dimensional control space.
2. The second is the local LSQ-to-Taylor result for algebraic discrete signed distance fields. On regular regions, this gives a mathematical justification for the practically important least-squares construction by showing that LSQ samples inherit the same local approximation order as Taylor samples.

3. The third is the subdivision result for manifold-valued data. By replacing affine averages with Karcher means, one obtains a companion multiresolution construction. This comes together with a continuous convergence theorem proved by proximity to the underlying linear scheme.

Taken together, these contributions support the thesis claim that higher-order signed distance representations, modular energy design, and manifold-aware optimization form a coherent mathematical basis for design-preserving surface retargeting.

7.2 Limitations

The presented method has several limitations. First, the main formulation is static: the target is treated as fixed, and the thesis does not attempt a full physical simulation of time-dependent elastic interaction. Second, the approximation theory for ADSDFs is local and applies on regular regions; it is not a global smoothness theory across the cut locus. In other words, there is no analysis of the failure modes of the LSQ technique across the cut locus. Third, the practical optimization procedures use damping, continuation, and related numerical safeguards whose success is motivated by experience and by existing optimization theory, but not analyzed here in full generality. Finally, the subdivision chapter is meant as a companion result rather than as a fully developed multiresolution framework: it proves continuous convergence for a stationary manifold-valued construction, but it does not address adaptive refinement, coarsening, or topology-changing updates.

7.3 Future directions

7.3.1 Physically motivated dynamic extensions

A natural next step is to replace some of the purely geometric constraints in the present framework with physically motivated energies while keeping the same basic target-side representation. Local rotation fields already play a central role in co-rotational and stiffness-warping models of elastic simulation [26, 27]. At the same time, SDF-based contact against rigid geometry has been developed in engineering and graphics settings, including corotational dynamics against rigid sur-

faces represented by discrete signed distance fields and local optimization for mesh–SDF contact generation [28, 29]. This suggests a concrete program: formulate implicit time-stepping problems in which the deformable body carries the same kind of rotation-augmented state used here, while the static obstacle is represented by an SDF, and then study whether higher-order SDF queries improve the stability and convergence of contact solves. This direction appears especially promising, but it also requires the most additional modeling choices, since constitutive laws, time integration, and contact mechanics would all have to be incorporated explicitly.

7.3.2 Adaptive and multilevel manifold-valued subdivision

The subdivision machinery used in this thesis is deliberately lightweight: it serves as a coarse-to-fine companion construction, but it is not yet a full multiresolution framework. Existing subdivision literature provides much stronger multilevel ideas, including interactive multiresolution editing, hierarchical refinement overlays, localized adaptive subdivision, and dynamic refinement or coarsening strategies [30–33]. A natural continuation is therefore to extend these techniques to the manifold-valued setting considered here. Two concrete questions stand out. The first is how to perform localized refinement, coarsening, or even connectivity-improving operations while updating quaternionic or manifold-valued data coherently. The second is how to optimize over a low-dimensional control mesh while evaluating the actual retargeting energies on the subdivided surface it generates. Because this direction builds directly on the subdivision chapter and on well-established Euclidean multiresolution techniques, it appears especially close at hand.

7.3.3 Adaptive and compressive higher-order SDFs

A second direction that appears especially natural is to move from uniform higher-order ADSDFs to adaptive and compressed constructions. There is already substantial literature on adaptive distance-field sampling and sparse hierarchical storage [34, 35], as well as on higher-order and *hp*-adaptive signed distance approximations [36–38]. What is missing, and what the present thesis makes newly plausible, is a representation that preserves the main ADSDF advantage—analytic low-order derivatives—while also adapting resolution to the needs of the optimizer. This suggests

refinement criteria driven not only by geometric approximation error, but also by optimization sensitivity: refine where contact terms are active, where Hessian information is needed, or where the LSQ-to-Taylor regime is most important to preserve. Among the future directions considered here, this one appears the most direct extension of the ADSDF theory developed in the thesis.

Taken together, these directions have different levels of immediacy. The dynamical extension appears highly promising, but it requires the introduction of constitutive modeling and time integration beyond the present thesis. By contrast, the multilevel subdivision and adaptive/compressive SDF directions seem especially likely to produce concrete progress in the near term, because they extend mathematical and representational structures that are already central to the current work.

Appendix A

Reference Material on Smooth Signed Distance

Functions

This appendix collects supplementary smooth SDF material that is useful as a reference but not essential to the main narrative of Chapter 2. In particular, the formulas in Section A.3 are intended to be implementation-ready test cases: they provide exact signed distance functions, gradients, and Hessians for a small set of common primitives.

A.1 Basic facts

For convenience, we restate the signed-distance convention from Definition 2.1.1 before listing implementation formulas. The restatement is included so that the exact primitive formulas in this appendix can be read independently of Chapter 2.

Definition A.1.1.

Let $\Omega \subset \mathbb{R}^3$ have nonempty boundary. The signed distance function of Ω is

$$\Phi_{\text{SDF}}(\mathbf{x}) = \begin{cases} -\text{dist}(\mathbf{x}, \partial\Omega), & \mathbf{x} \in \Omega, \\ \text{dist}(\mathbf{x}, \partial\Omega), & \mathbf{x} \notin \Omega. \end{cases}$$

Standard results imply that Φ_{SDF} is continuous and 1-Lipschitz. At a point with a unique closest point on $\partial\Omega$, the function is locally as smooth as the target surface; at points with multiple closest points, differentiability can fail. The singular set is commonly described as the cut locus [3].

Proposition A.1.1.

At every regular point of Φ_{SDF} ,

$$\|\nabla\Phi_{\text{SDF}}(\boldsymbol{x})\| = 1.$$

If

$$\mathcal{S}_c = \{y \in \mathbb{R}^3 : \Phi_{\text{SDF}}(y) = c\}, \quad c = \Phi_{\text{SDF}}(\boldsymbol{x}),$$

then $\nabla\Phi_{\text{SDF}}(\boldsymbol{x})$ is a unit normal to $\mathcal{S}_c$ at $\boldsymbol{x}$.

A.2 Characterization by Eikonal and closest-point properties

The unit-gradient condition alone does not characterize a signed distance function. The missing ingredient is a closest-point relation.

Definition A.2.1.

A function $\Phi : \mathbb{R}^3 \rightarrow \mathbb{R}$ is said to satisfy the *Eikonal property* if

$$\|\nabla\Phi(\boldsymbol{x})\| = 1$$

almost everywhere. It is said to satisfy the *closest-point property* if

$$\Phi(\boldsymbol{x} - \Phi(\boldsymbol{x})\nabla\Phi(\boldsymbol{x})) = 0$$

almost everywhere.

Theorem A.2.1.

Under the natural regularity hypotheses needed to interpret the preceding formulae, a function

$\Phi : \mathbb{R}^3 \rightarrow \mathbb{R}$ is a signed distance function if and only if it satisfies the Eikonal and closest-point properties [39].

The practical meaning of the closest-point formula is that, at a regular point,

$$\underbrace{\mathbf{x} - \Phi_{\text{SDF}}(\mathbf{x}) \nabla \Phi_{\text{SDF}}(\mathbf{x})}_{\text{closest point on the zero level set}}$$

recovers the nearest point on the target surface.

A.3 Exact formulas for common primitives

The following formulas are convenient for unit tests, finite-difference checks, and sanity checks on derivative code.

A.3.1 Sphere

For the unit sphere centered at the origin, let

$$\rho = \sqrt{x^2 + y^2 + z^2}.$$

Then

$$\Phi_{\text{SDF}}(x, y, z) = \rho - 1,$$

$$\nabla \Phi_{\text{SDF}}(x, y, z) = \frac{1}{\rho} \begin{bmatrix} x \\ y \\ z \end{bmatrix}, \quad \rho \neq 0,$$

and

$$\nabla^2 \Phi_{\text{SDF}}(x, y, z) = \frac{1}{\rho} I - \frac{1}{\rho^3} \begin{bmatrix} x \\ y \\ z \end{bmatrix} \begin{bmatrix} x & y & z \end{bmatrix}, \quad \rho \neq 0.$$

A.3.2 Torus

For a torus about the y -axis with major radius R and minor radius r , define

$$s = \sqrt{x^2 + z^2}, \quad a = s - R, \quad q = \sqrt{a^2 + y^2}.$$

Then

$$\Phi_{\text{SDF}}(x, y, z) = q - r.$$

Its gradient is

$$\nabla \Phi_{\text{SDF}}(x, y, z) = \frac{1}{q} \begin{bmatrix} a x/s \\ y \\ a z/s \end{bmatrix}, \quad s \neq 0, \quad q \neq 0.$$

Writing

$$J = \begin{bmatrix} x/s & 0 & z/s \\ 0 & 1 & 0 \end{bmatrix}, \quad u = \begin{bmatrix} a \\ y \end{bmatrix}, \quad H_s = \frac{1}{s} \begin{bmatrix} 1 - \frac{x^2}{s^2} & 0 & -\frac{xz}{s^2} \\ 0 & 0 & 0 \\ -\frac{xz}{s^2} & 0 & 1 - \frac{z^2}{s^2} \end{bmatrix},$$

a compact Hessian formula is

$$\nabla^2 \Phi_{\text{SDF}} = J^\top \left(\frac{1}{q} I_2 - \frac{1}{q^3} u u^\top \right) J + \frac{a}{q} H_s.$$

A.3.3 Oriented plane

Let $\vec{n} \in \mathbb{R}^3$ be a unit normal and $\mathbf{x}_0 \in \mathbb{R}^3$ a point on the plane. Then

$$\Phi_{\text{SDF}}(\mathbf{x}) = \vec{n} \cdot (\mathbf{x} - \mathbf{x}_0), \quad \nabla \Phi_{\text{SDF}}(\mathbf{x}) = \vec{n}, \quad \nabla^2 \Phi_{\text{SDF}}(\mathbf{x}) = 0.$$

A.3.4 Point and segment primitives

These are not globally signed distance functions by themselves, but they are the local distance primitives that reappear in featurewise analyses of mesh SDFs [2].

For a point $\mathbf{x}_0$,

$$\begin{aligned} d_{\text{pt}}(\mathbf{x}) &= \|\mathbf{x} - \mathbf{x}_0\|, \\ \nabla d_{\text{pt}}(\mathbf{x}) &= \frac{\mathbf{x} - \mathbf{x}_0}{\|\mathbf{x} - \mathbf{x}_0\|}, \\ \nabla^2 d_{\text{pt}}(\mathbf{x}) &= \frac{1}{\|\mathbf{x} - \mathbf{x}_0\|} I - \frac{(\mathbf{x} - \mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0)^\top}{\|\mathbf{x} - \mathbf{x}_0\|^3}. \end{aligned}$$

For a segment with endpoints $\mathbf{x}_0, \mathbf{x}_1$ and direction $\vec{\mathbf{e}} = \mathbf{x}_1 - \mathbf{x}_0$, let

$$t_*(\mathbf{x}) = \text{clamp}\left(\frac{(\mathbf{x} - \mathbf{x}_0) \cdot \vec{\mathbf{e}}}{\|\vec{\mathbf{e}}\|^2}, 0, 1\right).$$

Then the unsigned distance to the segment is

$$d_{\text{seg}}(\mathbf{x}) = \|\mathbf{x} - (\mathbf{x}_0 + t_*(\mathbf{x})\vec{\mathbf{e}})\|.$$

On subregions where the closest point lies in the interior of the segment, this reduces locally to the distance to an affine line; at the endpoints it reduces to the point-distance formula above.

A.4 Curvature of level sets from the Hessian

The Hessian of a signed distance function is most useful when interpreted level-set by level-set.

Proposition A.4.1.

Let $\mathbf{x}$ be a regular point of Φ_{SDF} , and let

$$\mathcal{S}_c = \{y \in \mathbb{R}^3 : \Phi_{\text{SDF}}(y) = c\}, \quad c = \Phi_{\text{SDF}}(\mathbf{x}).$$

Then $\nabla^2\Phi_{\text{SDF}}(\boldsymbol{x})\nabla\Phi_{\text{SDF}}(\boldsymbol{x}) = 0$. The other two eigenvalues of $\nabla^2\Phi_{\text{SDF}}(\boldsymbol{x})$ are the principal curvatures of $\mathcal{S}_c$ at $\boldsymbol{x}$. Consequently

$$\underbrace{\text{tr}(\nabla^2\Phi_{\text{SDF}}(\boldsymbol{x}))}_{\text{mean curvature}} = \kappa_1 + \kappa_2$$

and

$$\underbrace{\frac{1}{2} \left(\text{tr}(\nabla^2\Phi_{\text{SDF}}(\boldsymbol{x}))^2 - \text{tr}(\nabla^2\Phi_{\text{SDF}}(\boldsymbol{x})^2) \right)}_{\text{Gaussian curvature}} = \kappa_1\kappa_2.$$

Appendix B

Quaternions

The goal of this appendix on quaternions is to serve as a reasonable introduction/review and as a reference guide for the many quaternionic formulas and facts which are easy to forget. The first half covers purely algebraic facts about the quaternions and the second half covers some geometric properties of the set of unit quaternions viewed as a manifold equipped with a metric.

B.1 Algebra of quaternions

One can realize the quaternions $\mathbb{H}$ as the quotient $\mathbb{R}\langle i, j \rangle / (i^2 + 1, j^2 + 1, ij + ji)$ where we define $k = ij$.

Theorem B.1.1.

There are distinguished elements $i, j, k \in \mathbb{H}$ satisfying

$$i^2 = j^2 = k^2 = ijk = -1$$

and moreover $\{1, i, j, k\}$ is a basis of $\mathbb{H}$ as a 4-dimensional vector space. In other words, for all $q \in \mathbb{H}$ we can write

$$q = w + xi + yj + zk$$

for some unique choice of $w, x, y, z \in \mathbb{R}$.

In other words, the quaternions are something like a 4-dimensional version of the complex numbers.

Theorem B.1.2.

Using the identity $i^2 = j^2 = k^2 = ijk = -1$ and the properties of $\mathbb{H}$ as an algebra over the

real numbers one can prove

$$ij = k, \quad jk = i, \quad ki = j$$

and

$$ji = -k, \quad kj = -i, \quad ik = -j.$$

Notice $\mathbb{H}$ is not a commutative algebra since, for example, $ij \neq ji$. One might think the quaternions are anticommutative but this is also not true since, for example, $(i + 2j)(i + k) = -3 - 3j$ but $(i + k)(i + 2j) = -3 + 3j$.

The notions of conjugation and absolute value from the complex numbers can be extended to the quaternions, and this even gives us a formula for multiplicative inverses. It is useful to define dot products of quaternions along the way.

Definition B.1.1.

The **conjugate** q^* of $q = w + xi + yj + zk \in \mathbb{H}$ is defined

$$q^* = w - xi - yj - zk.$$

Given $q_1, q_2 \in \mathbb{H}$ with $q_1 = w_1 + x_1i + y_1j + z_1k$ and $q_2 = w_2 + x_2i + y_2j + z_2k$ we define the **dot product**

$$q_1 \cdot q_2 = w_1w_2 + x_1x_2 + y_1y_2 + z_1z_2.$$

The **absolute value** $|q|$ of $q = w + xi + yj + zk \in \mathbb{H}$ is defined

$$|q|^2 = w^2 + x^2 + y^2 + z^2 = qq^* = q^*q.$$

Theorem B.1.3.

Given $0 \neq q = w + xi + yj + zk \in \mathbb{H}$, the multiplicative inverse q^{-1} is given by

$$q^{-1} = \frac{q^*}{|q|^2} = \frac{w - xi - yj - zk}{w^2 + x^2 + y^2 + z^2}.$$

The next definitions and theorems connect quaternions to vector algebra.

Definition B.1.2.

Given $q = w + xi + yj + zk \in \mathbb{H}$ we call w the **real part** of the quaternion and $xi + yj + zk$ the **imaginary part** or **vector part** of the quaternion. If $w = 0$ we call q a **vector quaternion**.

We identify the set of vector quaternions with $\mathbb{R}^3$. In other words we will treat $(x \ y \ z)^T$ as equivalent to $xi + yj + zk$ so that we may mix quaternion operations with the usual vector operations in $\mathbb{R}^3$.

Theorem B.1.4.

Given $q_1 = w_1 + v_1, q_2 = w_2 + v_2 \in \mathbb{H}$ where $v_1 = x_1i + y_1j + z_1k$ and $v_2 = x_2i + y_2j + z_2k$

$$q_1 q_2 = (w_1 w_2 - v_1 \cdot v_2) + (w_1 v_2 + w_2 v_1 + v_1 \times v_2).$$

Finally we define the quaternionic logarithm and exponential.

Definition B.1.3.

We define $\exp: \mathbb{H} \rightarrow \mathbb{H}$, $\sin: \mathbb{H} \rightarrow \mathbb{H}$, and $\cos: \mathbb{H} \rightarrow \mathbb{H}$ by

$$\exp(q) = \sum_{n=0}^{\infty} \frac{q^n}{n!},$$

$$\sin(q) = \sum_{n=0}^{\infty} \frac{(-1)^n q^{2n+1}}{(2n+1)!},$$

$$\cos(q) = \sum_{n=0}^{\infty} \frac{(-1)^n q^{2n}}{(2n)!}.$$

Theorem B.1.5.

The quaternionic power series $\exp$, $\sin$, and $\cos$ converge (with respect to the norm $|\cdot|$) for all $q \in \mathbb{H}$. In fact, they converge absolutely in the sense that each of

$$\sum_{n=0}^{\infty} \frac{|q|^n}{n!},$$

$$\sum_{n=0}^{\infty} \frac{|q|^{2n+1}}{(2n+1)!},$$

$$\sum_{n=0}^{\infty} \frac{|q|^{2n}}{(2n)!},$$

converge for all $q \in \mathbb{H}$.

From these facts we can derive the following.

Theorem B.1.6.

Given $q = w + v \in \mathbb{H}$ where w is the real part of q and v is the vector part

$$\exp(q) = e^w \left(\cos(|v|) + \frac{v}{|v|} \sin(|v|) \right).$$

This inspires us to define the logarithm as follows.

Definition B.1.4.

We define $\log: \mathbb{H} \rightarrow \mathbb{H}$ by

$$\log(q) = \log(|q|) + \frac{v}{|v|} \arccos\left(\frac{w}{|q|}\right)$$

for $q = w + v \in \mathbb{H}$ where w is the real part of q and v is the vector part

B.2 Geometry of quaternions

The next definition and theorem establish the unit quaternions as a manifold.

Definition B.2.1.

A quaternion $q \in \mathbb{H}$ is called a **unit quaternion** in case $|q| = 1$.

We will identify the set of unit quaternions with S^3 , the 3-sphere. This choice is justified by the following theorem.

Theorem B.2.1.

The set $U = \{q \in \mathbb{H}: |q| = 1\}$, endowed with the subspace topology, is homeomorphic to the 3-sphere.

Famously, quaternions encode rotations in $\mathbb{R}^3$. The following theorem gives a formula.

Theorem B.2.2.

Let $v \in \mathbb{R}^3$, let $a \in \mathbb{R}^3$ be a unit vector, and let v' be the vector obtained by rotating v about the axis a through an angle $\theta \in \mathbb{R}$. Identifying vectors in $\mathbb{R}^3$ with pure imaginary quaternions, we have

$$v' = q_a v q_a^*$$

where

$$q_a = \cos(\theta/2) + a \sin(\theta/2) \in S^3.$$

The most natural way to represent rotations is with elements of $SO(3)$, that is, the manifold of 3×3 rotation matrices. The connection to quaternions is that the manifold S^3 *double covers* $SO(3)$. What this means in light of Theorem B.2.2 is that there are two distinct unit quaternions which represent a given rotation. In particular, we could have written $v' = (-q_a)v(-q_a)^*$. That is to say, given a unit quaternion q representing a rotation, the unit quaternion $-q$ represents the same rotation. This non-uniqueness of representation turns out not to be much of an obstruction in practical computations, but it should be kept in mind.

There is a natural choice³ of metric tensor g which gives S^3 the structure of a Riemannian manifold. This gives us a notion of distance $d_g: S^3 \times S^3 \rightarrow \mathbb{R}$ as well as exponential and logarithmic maps $\exp_p: T_p S^3 \rightarrow S^3$, $\log_p: S^3 \rightarrow T_p S^3$. The following theorems connect these to the quaternion operations we've discussed so far.

Theorem B.2.3.

The distance function $d_g: S^3 \times S^3 \rightarrow \mathbb{R}$ induced by the standard round metric on S^3 is given by

$$d_g(q_1, q_2) = |\log((q_2^*)q_1)| = \arccos(q_1 \cdot q_2).$$

If we think of S^3 as the unit 3-sphere embedded in $\mathbb{R}^4$, then this distance is the angle subtended by the great-circle arc connecting q_1 and q_2 .

Theorem B.2.4.

We identify each tangent space $T_p S^3$ with $\mathbb{R}^3$. Recall that we have already identified $\mathbb{R}^3$ with the set of vector quaternions. In other words, quaternions of the form $v = xi + yj + zk$ can be treated as elements of $T_p S^3$.

³at least, defined up to a constant

The exponential map $\exp_p: T_p S^3 \rightarrow S^3$ is given by

$$\exp_p(v) = p \exp(p^* v)$$

where the second $\exp$ is the exponential map on quaternions defined above.

Similarly, the logarithmic map $\log_p: S^3 \rightarrow T_p S^3$ is given by

$$\log_p(q) = p \log(p^* q)$$

where the second $\log$ is the logarithmic map on quaternions defined above.

Bibliography

- [1] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In Alexander Belyaev and Michael Garland, editors, *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*, pages 109–116. The Eurographics Association, 2007.
- [2] Gábor Valasek and Róbert Bán. Higher-order algebraic signed distance fields. *Computer-Aided Design and Applications*, 20(5):1005–1028, 2023.
- [3] David Gilbarg and Neil S. Trudinger. *Elliptic Partial Differential Equations of Second Order*. Classics in Mathematics. Springer, Berlin and Heidelberg, 2 edition, 2001.
- [4] Róbert Bán and Gábor Valasek. First-order signed distance fields. In *Eurographics 2020 Short Papers*, pages 33–36. The Eurographics Association, 2020.
- [5] John C. Hart. Sphere tracing: A geometric method for the antialiased ray tracing of implicit surfaces. *The Visual Computer*, 12(10):527–545, 1996.
- [6] Karsten Grove and Hermann Karcher. How to conjugate C^1 -close group actions. *Mathematische Zeitschrift*, 132:11–20, 1973.
- [7] Hermann Karcher. Riemannian center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics*, 30(5):509–541, 1977.
- [8] Knut Hüper and Jochen Trumpf. Newton-like methods for numerical optimization on manifolds. In *Conference Record of the Thirty-Eighth Asilomar Conference on Signals, Systems and Computers*, volume 1, pages 136–139. IEEE, 2004.
- [9] Edwin Catmull and James Clark. Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-Aided Design*, 10(6):350–355, 1978.
- [10] Charles Loop. Smooth subdivision surfaces based on triangles. Master’s thesis, Department of Mathematics, University of Utah, Salt Lake City, UT, 1987.

- [11] Denis Zorin and Peter Schröder. Subdivision for modeling and animation. ACM SIGGRAPH 2000 Course Notes, 2000.
- [12] Johannes Wallner and Nira Dyn. Convergence and C^1 analysis of subdivision schemes on manifolds by proximity. *Computer Aided Geometric Design*, 22(7):593–622, 2005.
- [13] Johannes Wallner. Smoothness analysis of subdivision schemes by proximity. *Constructive Approximation*, 24(3):289–318, 2006.
- [14] Andreas Weinmann. Nonlinear subdivision schemes on irregular meshes. *Constructive Approximation*, 31(3):395–415, 2010.
- [15] Annika Oehri, Philipp Herholz, and Olga Sorkine-Hornung. Higher-order continuity for smooth as-rigid-as-possible shape modeling. *Journal of Computer Graphics Techniques*, 14(1):198–215, 2025.
- [16] James A. Sethian. Fast marching methods. *SIAM Review*, 41(2):199–235, 1999.
- [17] Nicolas Boumal. *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, Cambridge, 2023.
- [18] Pierre-Antoine Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008.
- [19] Paul Breiding and Nick Vannieuwenhoven. Convergence analysis of Riemannian Gauss–Newton methods and its connection with the geometric condition number. *Applied Mathematics Letters*, 78:42–50, 2018.
- [20] Sho Adachi, Takayuki Okuno, and Akiko Takeda. Riemannian Levenberg–Marquardt method with global and local convergence properties, 2022.
- [21] Daniel W. H. Doo and Malcolm A. Sabin. Behaviour of recursive division surfaces near extraordinary points. *Computer-Aided Design*, 10(6):356–360, 1978.

- [22] Denis N. Zorin. *Stationary Subdivision and Multiresolution Surface Representations*. PhD thesis, California Institute of Technology, 1998.
- [23] Joe Warren and Henrik Weimer. *Subdivision Methods for Geometric Design: A Constructive Approach*. Morgan Kaufmann, San Francisco, CA, 2002.
- [24] Ulrich Reif. A unified approach to subdivision algorithms near extraordinary vertices. *Computer Aided Geometric Design*, 12(2):153–174, 1995.
- [25] Bijan Afsari. Riemannian L^p center of mass: Existence, uniqueness, and convexity. *Proceedings of the American Mathematical Society*, 139(2):655–673, 2011.
- [26] Matthias Müller, Julie Dorsey, Leonard McMillan, Robert Jagnow, and Barbara Cutler. Stable real-time deformations. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 49–54. Association for Computing Machinery, 2002.
- [27] Rahul Narain, Matthew Overby, and George E. Brown. ADMM $\supseteq$ projective dynamics: Fast simulation of general constitutive models. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 21–28. The Eurographics Association, 2016.
- [28] Miquel Aguirre and Stéphane Avril. An implicit 3D corotational formulation for frictional contact dynamics of beams against rigid surfaces using discrete signed distance fields. *Computer Methods in Applied Mechanics and Engineering*, 371:113275, 2020.
- [29] Miles Macklin, Kenny Erleben, Matthias Müller, Nuttapong Chentanez, Stefan Jeschke, and Zach Corse. Local optimization for robust signed distance field collision. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 3(1):8:1–8:17, 2020.
- [30] Denis Zorin, Peter Schröder, and Wim Sweldens. Interactive multiresolution mesh editing. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, pages 259–268. Association for Computing Machinery, 1997.

- [31] David R. Forsey and Richard H. Bartels. Hierarchical B-spline refinement. *ACM SIGGRAPH Computer Graphics*, 22(4):205–212, 1988.
- [32] Jun-Hai Yong and Fuhua Frank Cheng. Adaptive subdivision of Catmull–Clark subdivision surfaces. *Computer-Aided Design and Applications*, 2(1–4):253–261, 2005.
- [33] Zachary Ferguson, Teseo Schneider, Danny M. Kaufman, and Daniele Panozzo. In-timestep remeshing for contacting elastodynamics. *ACM Transactions on Graphics*, 42(4):105:1–105:15, 2023.
- [34] Sarah F. Frisken Gibson, Ronald N. Perry, Alyn P. Rockwood, and Thouis R. Jones. Adaptively sampled distance fields: A general representation of shape for computer graphics. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, pages 249–254. Association for Computing Machinery, 2000.
- [35] Ken Museth. VDB: High-resolution sparse volumes with dynamic topology. *ACM Transactions on Graphics*, 32(3):27:1–27:22, 2013.
- [36] Dan Koschier, Crispin Deul, and Jan Bender. Hierarchical hp-adaptive signed distance fields. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 189–198. The Eurographics Association, 2016.
- [37] Dan Koschier, Crispin Deul, Magnus Brand, and Jan Bender. An hp-adaptive discretization algorithm for signed distance field generation. *IEEE Transactions on Visualization and Computer Graphics*, 23(10):2208–2221, 2017.
- [38] Eduard Pujol and Antonio Chica. Adaptive approximation of signed distance fields through piecewise continuous interpolation. *Computers & Graphics*, 114:337–346, 2023.
- [39] Zoë Marschner, Silvia Sellán, Hsueh-Ti Derek Liu, and Alec Jacobson. Constructive solid geometry on neural signed distance fields. In *SIGGRAPH Asia 2023 Conference Papers*, pages 121:1–121:12. Association for Computing Machinery, 2023.