

Conformal Cone Parameterization through Optimal Control

Thesis by
Yousuf Soliman

Department of Mathematical Sciences
Carnegie Mellon University
Pittsburgh, Pennsylvania

Submitted in partial fulfillment of
the requirements for the degree of
Master of Science in Mathematical Sciences
May 2018



Acknowledgments

I would like to begin by thanking Keenan Crane for introducing me to discrete differential geometry and computer graphics, and for pointing me towards this fascinating topic. I would also like to thank both Keenan and Dejan Slepčev for advising me throughout the course of this project. They have both spent many hours discussing this work with me, and I deeply appreciate it. Much of this work will also appear in


Yousuf Soliman, Dejan Slepčev, Keenan Crane
Optimal Cone Singularities for Conformal Flattening
ACM Transactions on Graphics 38 (4) 2018 (conditionally accepted)

and so I would also like to thank them for helping me write and revise this manuscript. I would be remiss not to thank all of the wonderful faculty here at CMU that helped me develop a solid foundational and technical background in analysis. In particular, I would like to thank Giovanni Leoni, Ian Tice, Irene Fonseca, and Dmitry Kramkov for teaching exceptional courses. Finally, I am grateful to Professor Hayden Schaeffer for serving on my thesis committee.

Abstract

Surface parameterizations with low metric distortion are essential for a wide variety of applications ranging from geometry processing, to digital manufacturing, to machine learning. Conformal parameterizations are easy to compute and exactly preserve angles, but can significantly distort areas. Cone singularities have been introduced as a way of mitigating area distortion, but finding the configuration of cones that best reduces distortion is notoriously difficult. This thesis develops a simple strategy that provides globally optimal configurations of singularities, in the sense that the configuration minimizes the total area distortion among all possible conformal cone configurations (number, placement, angle) that have no more than a fixed total cone angle. In practice, our optimal cone configurations can yield dramatically lower area distortion than those found via existing heuristics. Moreover, the approach can be extended to allow user-defined notions of importance, find the best flattening with a convex or polygonal boundary, and produce solutions with only positive cone angles.

Our approach can be summarized as follows: the cone singularity placement problem is relaxed to a convex optimization problem over the space of finite signed Radon measures with a sparsity inducing regularization. By utilizing Fenchel-Rockafellar duality we obtain an equivalent formulation over some usual function spaces, which are easily and properly discretized using finite elements. Computing the optimal configuration of cones amounts to solving a sequence of sparse linear systems easily built from the usual cotangent Laplacian. The experimental results presented in this thesis provide some evidence that simple sparsity inducing norms may be widely applicable to problems arising in computer graphics and geometry processing.



Left: Conformal flattening yields maps with no angle distortion, but can exhibit severe area distortion. *Right:* An optimal arrangement of a small number of cone singularities (here, just nine) can push a conformal flattening very close to perfect isometry.

Contents

Acknowledgments	iv
Abstract	v
1 Introduction	1
2 Background	6
2.1 Linear and Convex Duality	6
2.2 Measure Theory	8
2.3 Riemannian Geometry	10
2.4 Sobolev Spaces on Manifolds	11
3 Problem Setup	13
3.1 Cone Singularities and the Yamabe Equation	13
3.2 Basic Problem	15
3.3 Measures of Area Distortion	17
3.4 Relaxation	18
3.5 Generalizations	20
4 Theoretical Analysis	22
4.1 Abstract Framework	22
4.2 Poisson Equation with Measures	24
4.3 Existence and Duality for Optimal Cones	28
4.4 Regularity and Rounding	35
4.5 Moreau-Yosida Regularization	41
5 Discretization	48
5.1 Preliminaries	48
5.2 Semismooth Newton Method	49
5.3 Algorithm	53
6 Results	60
6.1 Validation and Comparisons	60
6.2 Extensions	64
7 Discussion and Conclusions	67
References	69

Chapter 1

Introduction



Figure 1.1: A conformal cone parameterization is equivalent to flattening a smooth surface, like the sphere, over a polyhedral surface, which can be cut and unfolded into the plane without further distortion. By adding more and more cone points, one can make area distortion arbitrarily small. (Texture courtesy NASA Earth Observatory.)

Mesh parameterization is a fundamental component of a wide variety of problems in applied geometry—beyond traditional tasks in computer graphics (such as texture mapping), surface flattenings have become an important component in a diverse collection of areas ranging from digital manufacturing to machine learning [41, 44]. Ideally, one would like a parameterization that is *isometric*, *i.e.*, no distortion of lengths or areas, but for general curved surfaces no such map exists. Conformal flattenings are attractive because they completely eliminate angle distortion, and are easily computed via linear or convex problems. However, they can also yield significant distortion of areas, which is problematic for applications since a large region of the surface is represented by only a very tiny region in the parameter domain.

The basic idea behind *cone flattening* [39] is that, intrinsically, many surfaces look more like a polyhedron than the flat plane—consider for instance maps of the Earth generated by conformally mapping the globe onto a regular polyhedron (Fig. 1.1). Since this initial map induces very little area distortion, and since the polyhedron can then be cut and unfolded into the plane without further stretching, the composite map also has low area distortion. Of course, different polyhedral metrics will lead to different amounts of distortion—the problem of cone parameterization therefore boils down to deciding on a configuration of vertices, which is described by the number, placement, and angle of the associated *cone singularities*. As stated, however, this problem is ill-posed: one can always reduce distortion further by considering a finer polyhedron, *i.e.*, allowing more cones. To make this problem well-posed, one can ask for the best configuration with a fixed number of cones, or alternatively (as we will do), fix the total magnitude $\Phi := \sum_i |\phi_i|$ of all cone angles ϕ_i .

Though a number of strategies have been developed for picking cones, none come with a clear guarantee of optimality, and in practice each can be confounded by certain types of models. The strat-

egy we develop ensures that total area distortion is always globally minimized, thereby providing substantial robustness and, in practice, significantly lower distortion on many real-world examples. The key insight is that this seemingly hard combinatorial problem can be relaxed to an easier convex problem, whose minimizers can still be used to obtain optimal solutions. In particular, we formulate a convex optimization problem over the space of finite signed Radon measures, and we add a sparsity inducing regularization to produce measures that represent cone singularities. By using convex duality, we obtain an equivalent optimization problem framed over the space of continuous functions and a set of first order optimality conditions. The resulting optimality system can be efficiently solved numerically using semismooth Newton methods (Chapter 5). In the end, we obtain a practical and efficient algorithm that:

- finds conformal flattenings of minimal total area distortion among all possible configurations of cones and choices of boundary conditions,
- can be trivially accelerated with a simple multi-resolution scheme,
- provides user control over both regions where cones can be placed, as well as regions where distortion should be measured,
- provides the ability to find optimal flattenings with a convex or polygonal boundary, and
- allows cone angles to be limited to a given range (*e.g.*, positive only or $[-\pi/2, \pi/2]$).

Beyond simply developing an algorithm, we also start to develop an understanding of some fundamental questions which can help to inform algorithmic decisions both in the present thesis and in the development of future work. In particular, we look at the practical importance of choosing a principled measure of area distortion (Sec. 3.3), we analyze stability of cone flattenings with respect to perturbations of the cones, *i.e.*, how much will distortion change if singularities are “merged” or “split” (Sec. 4.4); and we consider the approximability of smooth metrics by polyhedral cone metrics from an analytical point of view, *i.e.*, when can a given metric be arbitrarily well-approximated by cones (Sec. 4.4). We also carefully analyze the solutions to our optimization problem—a subtle point is that, in the continuous setting, minimizers of the relaxed problem live in the measure space H^{-1} and hence cannot exactly describe cones, which correspond to Dirac delta measures. In practice we can therefore (very rarely) get tiny clusters of cones; the stability result mentioned above ensures that these clusters can be rounded to a nearby cone configuration with a virtually imperceptible change in area distortion (Fig. 4.1).

Related Work

A variety of problems in digital geometry processing seek to determine ideal locations for certain singular features. For instance, in the design of tangent vector fields, tensor fields or rotationally symmetric direction fields, judicious placement of singularities can have a significant effect on the global regularity of the field [65]. These types of singularities are tangentially related to the ones we consider here, in the sense that such fields are often used to drive surface flattening [11], though the link is fairly indirect: singularities that yield highly smooth fields do not immediately guarantee a good flattening (especially when far from integrable). In this setting one can find singularities that yield optimal

smoothness by simply solving a sparse eigenvalue problem [40]. There is also an extensive literature on *cutting* surfaces into pieces that can be flattened with low area distortion [51]; an important distinction between general cuts and those arising from a cone flattening is that the latter is automatically *seamless* away from a small isolated set of cone points—consider for instance “painting” on the intrinsic polyhedral domain (Fig. 1.1, *top*) rather than the final 2D layout (Fig. 1.1, *bottom*). To date there is also no method for surface cutting that guarantees global optimality in the sense of area distortion (*e.g.*, subject to a bound on the length of the cut); in fact even the simpler problem of finding the *shortest* way of cutting a surface into a disk is NP-hard [24].

Conformal Cone Singularity Placement In this thesis, we focus specifically on placing *cone singularities* to reduce area distortion in conformal flattening; here a variety of strategies have been proposed. Kharevych et al. [39] initially investigated cone flattenings by manually drawing layouts and adjusting cone angles to reduce distortion. Springborn et al. [54] propose a method for cone flattening (**CETM**) where cones are chosen via a simple greedy algorithm: iteratively flatten the mesh; at each iteration place a new cone singularity at the point of greatest area distortion. In subsequent iterations, cone points are effectively treated as punctures in the domain, leading to cone angles that are automatically determined by the flattening process. As we will discuss in Sec. 3.3, however, this approach is mesh dependent since the Dirichlet energy of the log conformal factor blows up in the presence of cone singularities. In a parallel development, Ben-Chen et al. [8] devise an algorithm for cone flattening (**CPMS**) where cone locations are chosen via the same greedy strategy, but angles are instead determined via a diffusion process involving Gaussian curvature; this basic strategy was recently accelerated by Vintescu et al. [68] using hierarchical persistence. Later, Myles & Zorin [45] develop a method for seamless global parameterization (**GPIF**), where the first stage is to determine the cone configuration by incrementally flattening the surface starting with the flattest regions, *i.e.*, those with smallest Gaussian curvature. A key insight of our work is that curvature does not always provide useful information about how cones should be arranged, since such reasoning does not account for the cones’ non-local influence on area distortion. In fact, one can find many examples where the optimal configuration places cones in regions that are flat—see for instance Fig. 5.2 where an optimal cone configuration includes cones in intrinsically flat regions; another example of how curvature-based approaches may lead to suboptimal solutions is shown in Fig. 1.2. We instead adopt a different point of view, namely that the problem of finding optimal cones can be better understood as an approximation problem—for instance, if the surface were first flattened without cone singularities, one should seek the best approximation of the resulting log conformal factor u by a finite sum of harmonic Green’s functions, which do not have compact support. A rather surprising result is that for this problem a sparsity-inducing norm will concentrate the solution onto isolated *points* (*i.e.*, cones) rather than curves or other regions.

Cone Metrics, Orbifolds, and Liouville Equations More broadly, Riemannian metrics with conical singularities can be understood from a variety of different points of view. Thurston studied a geometric picture of Riemannian metrics with an *orbifold structure*, *i.e.*, each point must locally look Euclidean or like the quotient of a Euclidean space under the action of a discrete group [56, Chapter 13]; he also showed that every cone metric can be triangulated, providing a clear connection between cone manifolds and the triangle meshes used in computer graphics [57]. Very recently this orbifold perspective has become quite fruitful in digital geometry processing, leading to a variety of algorithms for computing canonical mappings between surfaces with landmarks [61], and surface parameterization

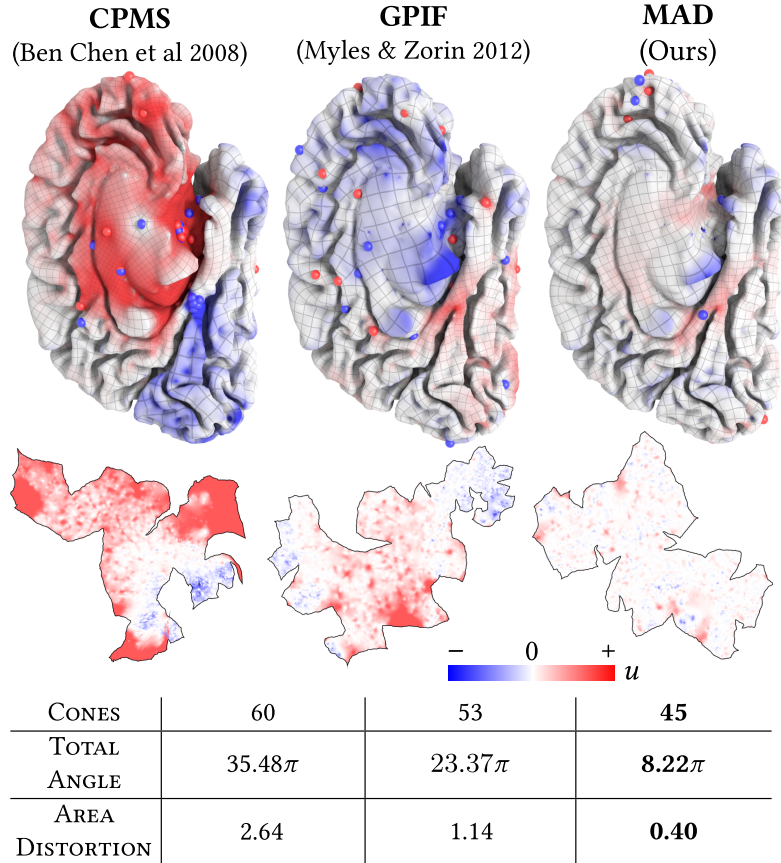


Figure 1.2: Even with fewer cones and much smaller total cone angle, our cone placement strategy (MAD) yields far lower area distortion than previous methods. This effect is especially apparent on shapes like the brain, which do not have obvious peaks of curvature.

algorithms with guarantees on global injectivity [2, 3, 4]. A very different perspective on cone metrics centers around the more analytical point of view of *Liouville equations*, in particular the Yamabe equation $\Delta u = K_0 - e^{2u}K$ describing the change in Gaussian curvature K under a pointwise conformal scaling $g = e^{2u}g_0$ of a metric g_0 . Troyanov provided some of the early foundations for studying this equation in the context of singular cone metrics [59, 60], which continues to be studied [20, 22, 23]. This intrinsic, analytic point of view has been used as a starting point for many recent algorithms in computational conformal geometry including the ones mentioned above (CPMS, CETM, and GPIF), as well as a recent method for conformal flattening [49]; it also plays a fundamental role in the method we develop here.

Convex Optimization and Semismooth Operator Equations Carefully formulating the cone placement problem as a PDE constrained optimization problem allows us to take advantage of highly effective optimization methods that have been recently developed by the optimal control community for *semismooth operator equations*. A key insight of these methods is that one should not directly discretize the original problem, but first formulate optimality conditions involving both primal and dual variables in the continuous setting, then discretize these optimality conditions [28, 36, 52]. The reason is that

directly discretizing the original problem may destroy important structures and relationships that appear in the context of continuous function (or measure) spaces, but are lost when moving to finite dimensional discrete spaces. A primal/dual formulation of our problem yields optimality conditions described by a semismooth operator equation, which can be solved effectively using *semismooth Newton methods* [16, 62, 64]. In particular, we mainly follow the approach of Hinze [34], where one does not need to directly optimize the control variables (in our case, the curvature measure used as a proxy for cone singularities) but instead introduces a collection of *adjoint variables*, which in our case amount to the Laplace *inverse* of the log conformal scale factors. These variables are highly unusual and do not appear in previous work on conformal flattening. For problems involving highly irregular solutions (like our cone distribution) it is also important to properly regularize this problem—here we apply *Moreau-Yosida regularization* to counteract numerical instability and improve the rate of convergence [9, 32, 35, 37, 58]. Finally, since we solve a relaxed problem we need a way to encourage sparsity; recent work by Clason and others provides a rigorous foundation for applying sparsity-inducing *measure norms* to PDE constrained optimization problems, mirroring how ℓ^1 norms are used to encourage sparsity for purely discrete problems [13, 18, 19]. Our work builds on all of this literature and shows how a similar formulation can be applied to problems in geometry processing and computer graphics.

Chapter 2

Background

In this chapter we review the main notions from functional analysis, measure theory, convex analysis, and differential geometry that will be used throughout our work.

2.1 Linear and Convex Duality

2.1.1 Functional Analysis

The *dual* of a normed vector space $(V, \|\cdot\|_V)$ is the space of linear continuous functionals $V \rightarrow \mathbb{R}$, and is denoted by V^* . We endow V^* with the *dual norm*

$$\|L\|_{V^*} := \sup \{L(v) : v \in V, \|v\|_V \leq 1\},$$

which makes V^* into a Banach space. The notation $\langle L, v \rangle_{V^* \times V} := L(v)$, denotes the duality pairing between V^* and V . Often, we will just write $\langle \cdot, \cdot \rangle = \langle \cdot, \cdot \rangle_{V^* \times V}$ if the vector spaces in question are clear from context. This notation is nice since many of the intuitions we have about inner products in Hilbert spaces carry over to arbitrary normed vector spaces once we replace the inner product with a duality pairing. Recall that a linear operator $T : V \rightarrow W$ is continuous if and only if $\|T\|_{\mathcal{L}(V, W)} < +\infty$.

Let $\Lambda : V \rightarrow W$ be a continuous linear operator between normed vector spaces. The *adjoint* is the map $\Lambda^* : W^* \rightarrow V^*$ characterized by

$$\langle L, \Lambda v \rangle_{W^* \times W} = \langle \Lambda^* L, v \rangle_{V^* \times V}$$

for all $L \in W^*$, $v \in V$.

We can use the duality pairing between V^* and V to endow V^* with the weak- * topology, denoted by $\sigma(V^*, V)$. This is the topology generated by finite intersections of sets of the form

$$\left\{ L \in V^* : |L(v)| < \frac{1}{n} \right\},$$

where $v \in V$ and $n \in \mathbb{N}$. It is a classical result that $\sigma(V^*, V)$ makes V^* into a locally convex topological vector space. We write $L_n \xrightarrow{*} L$ to indicate that L_n converges to L in the weak- * topology. An equivalent characterization of weak- * convergence of functionals is

$$L_n \xrightarrow{*} L \iff L_n(v) \rightarrow L(v) \quad \text{for all } v \in V$$

as $n \rightarrow +\infty$. A central appeal of the weak- * topology is that the unit ball is compact:

Theorem 1 (Banach-Alaoglu). *Let V be a normed vector space. Then the unit ball*

$$B_{V^*}(\mathbf{0}, 1) = \{L \in V^* : \|L\|_{V^*} \leq 1\}$$

*is compact in the weak- * topology.*

Finally, the *convex conjugate* of any function $F : X \rightarrow \overline{\mathbb{R}}$ is the function $F^* : X^* \rightarrow \overline{\mathbb{R}}$ given by

$$F^*(L) := \sup_{x \in X} \{L(x) - F(x)\}$$

A concrete example is that the convex conjugate of any squared norm is the corresponding (squared) *dual norm*—for example, the dual L^1 is L^∞ ; the dual of L^2 is just L^2 (making appropriate identifications). An example important in our setting is that the measure norm $\|\cdot\|_{\mathcal{M}}$ can be obtained as the convex conjugate of an indicator function on the unit ball of continuous functions.

2.1.2 Convex Analysis

A powerful tool in optimization is formulation of a *dual problem*, which may be easier to work with than the original *primal* problem. A very general purpose approach is *Lagrange duality*, though for problems involving measure spaces this approach becomes quite technical (see Sec. 4.3). We instead use the more specialized technique of *Fenchel-Rockafellar duality*, which for problems of the kind considered in this thesis easily yields an explicit characterization of minimizers.

In particular, suppose we want to solve the problem

$$\inf_{x \in X} F(x) + G(\Lambda x),$$

where (subject to mild technical conditions) $F : X \rightarrow \overline{\mathbb{R}}$ and $G : Y \rightarrow \overline{\mathbb{R}}$ are convex functions on normed vector spaces X and Y (*resp.*), and $\Lambda : X \rightarrow Y$ is a linear map. The Fenchel-Rockafellar duality theorem states that this problem is equivalent to the dual problem

$$\max_{y^* \in Y^*} -F^*(\Lambda^* y^*) - G^*(-y^*),$$

i.e., both problems have the same optimal value, and moreover, optimal points \bar{x} and \bar{y}^* can be related by an explicit set of optimality conditions, as described below.

Theorem 2 (Fenchel-Rockafellar duality). *Let V and W be Banach spaces, and let $\mathcal{F} : V \rightarrow \overline{\mathbb{R}}$ and $\mathcal{G} : W \rightarrow \overline{\mathbb{R}}$ be lower-semicontinuous, convex, and proper functions. Let $\Lambda : V \rightarrow W$ be a linear continuous operator. If there exists a point $v_0 \in V$ such that $\mathcal{F}(v_0) < +\infty$ and $\mathcal{G}(\Lambda v_0) < +\infty$, and \mathcal{G} is continuous at Λv_0 , then*

$$\inf_{v \in V} \{\mathcal{F}(v) + \mathcal{G}(\Lambda v)\} = \max_{w^* \in W^*} \{-\mathcal{F}^*(\Lambda^* w^*) - \mathcal{G}^*(-w^*)\}. \quad (2.1)$$

The equality in Eqn. 2.1 is attained by some pair (\bar{v}, \bar{w}^*) if and only if the following extremality conditions hold:

$$\begin{cases} \Lambda^* \bar{w}^* \in \partial \mathcal{F}(\bar{v}) \\ -\bar{w}^* \in \partial \mathcal{G}(\Lambda \bar{v}) \end{cases}$$

Optimality Conditions As with Lagrange duality, optimal points can be nicely characterized in terms of both primal and dual variables. For a differentiable objective, these conditions would simply involve derivatives of F and G . One challenge, however, is that these functions are not required to be differentiable (as will be the case in our problem, due to use of the measure norm—see Sec. 3.4). We therefore formulate optimality conditions in terms of the *subdifferential*. Intuitively, if the gradient provides the best linear approximation at a point, then the subdifferential describes all linear approximations “below” the function (see inset figure). More precisely, for any convex function f from a normed vector

space X to \mathbb{R} , the subdifferential is defined as

$$\partial f(x) := \{L \in X^* : L(y) - L(x) \leq f(y) - f(x) \text{ for all } y \in X\}.$$

A simple example is the function $f : \mathbb{R} \rightarrow \mathbb{R}$ given by $x \mapsto |x| + x^2$: for $x \neq 0$ the subdifferential contains just the ordinary derivative $f'(x)$; at $x = 0$, the subdifferential is $\partial f(0) = [-1, 1]$ (see inset figure). Using the subdifferential, we can express the optimality conditions as

$$\begin{cases} \Lambda^* \bar{y}^* & \in \partial F(\bar{x}), \\ -\bar{y}^* & \in \partial G(\Lambda \bar{x}). \end{cases}$$

For example, when Λ is the identity, and both F and G are differentiable, these conditions amount to saying that $\nabla F(\bar{x}) = -\nabla G(\bar{x})$, *i.e.*, the usual statement about Lagrange multipliers.

From the perspective of computation, a somewhat surprising outcome is that for problems involving measures, discretizing and solving these optimality conditions appears to yield numerical behavior far superior to simply trying to solve the primal or dual problem directly (Sec. 3.4). In other words, deriving the optimality conditions in the continuous setting and *then* discretizing is not equivalent to discretizing the optimization problem and then computing optimal solutions—the former approach seems to preserve essential structure from the continuous setting (namely, relationships between primal and dual variables).

2.2 Measure Theory

A key feature of our approach is that the underlying optimization problem is framed over the space of measures, instead of usual function spaces. Consider a compact topological space, Y . A *measure* on Y is a countably additive set function $\mu : \mathcal{B}(Y) \rightarrow \mathbb{R}$ defined on the Borel σ -algebra satisfying $\mu(\emptyset) = 0$. A Radon measure on a σ -finite topological space is a Borel measure that is finite on compact sets. We let $\mathcal{M}(Y)$ denote the space of all bounded Radon measures on Y . For $\mu \in \mathcal{M}(Y)$ we consider the *total variation measure* $|\mu| : \mathcal{B}(Y) \rightarrow \mathbb{R}^+$ given by

$$|\mu|(E) := \sup \left\{ \sum_{n=1}^N |\mu(E_n)| : E_n \in \mathcal{B}(Y), \bigcup_{n=1}^N E_n \subseteq E \right\}.$$

We also consider the space of *positive* measures, denoted by $\mathcal{M}^+(Y)$. Note that $|\mu| \in \mathcal{M}^+(Y)$. We can make $\mathcal{M}(Y)$ into a normed vector space by considering the total variation norm:

$$\|\mu\|_{\mathcal{M}(Y)} := |\mu|(Y).$$

In fact, $\mathcal{M}(Y)$ is a Banach space under this norm (this will follow from the characterization of $\mathcal{M}(Y)$ using duality).

This definition can be unweildly for purposes of optimization, and so we employ linear duality to obtain a nicer characterization of $\mathcal{M}(Y)$. Consider the vector space of all continuous functions, denoted by $\mathcal{C}(Y)$. We make $\mathcal{C}(Y)$ into a Banach space by endowing it with the norm $\|\varphi\|_{\mathcal{C}(Y)} := \sup_{x \in Y} |\varphi(x)|$. The Riesz representation theorem for $\mathcal{C}(Y)$ states that for every linear and continuous functional $L : \mathcal{C}(Y) \rightarrow \mathbb{R}$ there exists a unique Radon measure $\mu \in \mathcal{M}(Y)$ such that

$$L(\varphi) = \int_Y \varphi \, d\mu$$

for all $\varphi \in \mathcal{C}(Y)$ [25, Theorem 7.2]. Similarly, every $\mu \in \mathcal{M}(Y)$ defines a linear functional on $\mathcal{C}(Y)$ by

Lebesgue integration, as above. In light of this, we can make the identification $\mathcal{C}(Y)^* = \mathcal{M}(Y)$, and make $\mathcal{M}(Y)$ into a Banach space with the norm

$$\|\mu\|_{\mathcal{M}(Y)} = \sup \left\{ \int_Y \varphi \, d\mu : \varphi \in \mathcal{C}(Y), \|\varphi\|_{\mathcal{C}(Y)} \leq 1 \right\}.$$

The *Jordan decomposition theorem* allows us to uniquely write any $\mu \in \mathcal{M}(M)$ as:

$$\mu = \mu^+ - \mu^-,$$

where $\mu^\pm \in \mathcal{M}^+(Y)$ are positive measures. We can use the Jordan decomposition to express the total variation norm:

$$\|\mu\|_{\mathcal{M}(Y)} = |\mu|(Y) = \mu^+(Y) + \mu^-(Y).$$

Furthermore, note that $|\mu|$ is absolutely continuous with respect to μ , and so by the Radon-Nikodym theorem we have that

$$\frac{d\mu}{d|\mu|} \in L^1(Y, d|\mu|),$$

Furthermore, we have that $|d\mu/d|\mu|| = 1$ for $|\mu|$ -almost every $x \in Y$.

The main motivation for considering the space $\mathcal{M}(Y)$ is that the above measure norm is the natural generalization of the L^1 -norm on functions, which is known to promote sparsity. To see this, fix a measure ν on Y , and consider any integrable function $f \in L^1(Y, \nu)$. If we consider the measure

$$\mu_f(E) := \int_E f \, d\nu,$$

then we have $\|\mu_f\|_{\mathcal{M}(Y)} = \|f\|_{L^1(Y, \nu)}$. However, *not* every measure can be represented as a function in the above way. The prototypical example of this, and the most important example for our work, is the *Dirac delta measure*. For any fixed $x \in Y$ we define the Dirac delta measure supported at x to be

$$\delta_x(E) := \begin{cases} 1 & \text{if } x \in E, \\ 0 & \text{if } x \notin E. \end{cases}$$

This corresponds to the linear functional $L(\varphi) := \varphi(x)$.

Since $\mathcal{M}(Y) \cong \mathcal{C}(Y)^*$ we know that the unit ball in $\mathcal{M}(Y)$ is weak-* compact (that is, compact with respect to $\sigma(\mathcal{M}(Y), \mathcal{C}(Y))$). Compare this with the (weak) precompactness characterization of sets in $L^1(Y; d\nu)$:

Theorem 3 (Dunford-Pettis). *Let $\mathcal{F} \subset L^1(Y)$. The \mathcal{F} is weakly sequentially precompact if and only if*

- (i) \mathcal{F} is bounded in $L^1(Y)$,
- (ii) \mathcal{F} is equi-integrable and for every $\varepsilon > 0$ there exists $E \subset Y$ with $E \in \mathcal{B}(Y)$ such that $\mu(E) < +\infty$ and

$$\sup_{f \in \mathcal{F}} \int_{Y \setminus E} |f| \, d\nu \leq \varepsilon.$$

This equivalence between weak sequential precompactness in $L^1(Y)$ and equi-integrability (defined below) shows that the unit ball in $L^1(Y)$ is not weakly compact. This often makes L^1 an undesirable space to work in for the calculus of variations.

Definition 4. $\mathcal{F} \subseteq L^1(Y, d\nu)$ is *equi-integrable* if for every $\varepsilon > 0$ there exists $\delta > 0$ such that

$$\int_E |f| d\nu \leq \varepsilon$$

for all $f \in \mathcal{F}$ and all measurable $E \subset Y$ with $\nu(E) \leq \delta$.

2.3 Riemannian Geometry

In this section we very quickly recall the main objects of study in differential and Riemannian geometry, building up the necessary machinery to define the Laplace-Beltrami operator. A more complete exposition can be found in any standard textbook on Riemannian geometry [38].

Definition 5. A Riemannian metric on M is a section of the tensor bundle $T_0^2M = TM^* \otimes TM^*$

$$\begin{aligned} g &: M \rightarrow TM^* \otimes TM^*, \\ p &\mapsto g_p \in (T_pM)^* \otimes (T_pM)^* \cong (T_pM \times T_pM)^*, \end{aligned}$$

such that the bilinear map $g_p : T_pM \times T_pM \rightarrow \mathbb{R}$ is symmetric and positive definite.

Now we can define a *Riemannian manifold* as a pair (M, g) where M is a smooth manifold and g is a Riemannian metric on g . The Riemannian manifold will be the central object of study in our work.

Definition 6. The Riemannian volume form is the unique volume form $dA_g \in \Omega^n(M)$ satisfying

$$dA_g(e_1, \dots, e_n) = 1,$$

whenever (e_1, \dots, e_n) is a positive orthonormal basis of T_pM .

By an abuse of notation, we write for $f \in \mathcal{C}^\infty(M)$

$$\int_M f := \int_M f dA_g.$$

Often we will write $dA = dA_g$ when the metric is understood from context. The existence of a Riemannian volume form is equivalent to the orientability of the manifold.

Definition 7. The Laplace-Beltrami operator on (M, g) is the differential operator

$$\Delta_g := \star d \star d = \delta d$$

Proposition 8. In local coordinates, the Laplace-Beltrami operator is written as

$$\Delta_g = -\frac{1}{\sqrt{\det(g)}} \frac{\partial}{\partial x^j} \left(\sqrt{\det(g)} g^{ij} \frac{\partial f}{\partial x^i} \right).$$

Note that in local coordinates Δ_g is an elliptic differential operator in divergence form. Note that we could have alternatively just skipped all of the machinery above, but it is very instructive to use the machinery of the exterior algebra to define this since it provides a very natural approach to discretization. In particular, one simply needs to define discrete Hodge star and exterior derivative operators to define the Laplace-Beltrami operator. This is in fact the approach we take (which just happens to coincide with the finite element discretization with piecewise linear finite elements).

Now we briefly define an affine connection (also known as a covariant derivative) to be able to invariantly define the Sobolev spaces on manifolds.

Definition 9. A connection on a vector bundle $E \rightarrow M$ is an \mathbb{R} -linear map

$$\nabla : \Omega^0(M; E) \rightarrow \Omega^1(M; E)$$

such that if $f : M \rightarrow \mathbb{R}$ and $\sigma \in \Omega^0(M; E)$ we have the Leibniz product rule

$$\nabla(f\sigma) = df \otimes \sigma + f \nabla \sigma.$$

2.4 Sobolev Spaces on Manifolds

In this section we define the Sobolev spaces on Riemannian manifolds, recall the basic embedding and trace theorems, and prove a classical elliptic regularity theorem related to the Laplace-Beltrami operator.

Consider an *oriented* Riemannian manifold (M, g) . Recall that the Riemannian metric defines the Levi-Civita connection ∇^g and a Riemannian volume form $dA_g = \star 1$. Using the Riesz-Representation theorem for the space of measures, we deduce that dA_g defines a Radon measure on M .

Definition 10. The space $L^p(M, dA_g)$ is the space of (equivalence classes of) p -integrable functions on the measurable space $(M, \mathcal{B}(M), dA_g)$. When the metric or volume form is understood from context we write $L^p(M)$ for $L^p(M, dA_g)$.

Similarly, we can define the L^p -spaces for sections of a vector bundle $\pi : E \rightarrow M$ with a metric h . For simplicity, we assume that the fibers of E are separable so that notions of strong/Bochner measurability coincide with Borel measurability—this seems like a delicate point, but this shouldn't show up in our work since we will only consider the finite-dimensional cases of $E = TM^{\otimes k}$.

Definition 11. The space $L^p(M; E)$ is the space of equivalence classes of Borel measurable maps $\sigma : M \rightarrow E$ such that $(\pi \circ \sigma)(x) = x$ for dA_g -almost every $x \in M$ and such that the map $x \mapsto \|\sigma(x)\|_h$ is in $L^p(M, dA_g)$.

Now we can define the Sobolev spaces on manifolds in an invariant way quite easily!

Definition 12. The Sobolev space $W^{1,p}(M)$ is defined as the set of functions $u \in L^p(M; dA_g)$ such that there exists $v \in L^p(M; T^*M)$ satisfying

$$\int_M \langle v, \varphi \rangle dA_g = \int_M \langle u, \nabla^* \varphi \rangle dA_g$$

for all $\varphi \in \mathcal{C}_c^\infty(T^*M)$. Here $\nabla : \Omega^0(M) \rightarrow \Omega^1(M)$ is the Levi-Civita connection and ∇^* is the L^2 -adjoint of the connection.

Note that we can use this same definition to define higher order Sobolev spaces by replacing the connection by the iterated connection $\nabla^k : \Omega^0(M) \rightarrow \Omega^k(M)$. Although this definition is quite nice to write down, it is often more convenient (or necessary) to work in local coordinates.

Proposition 13. Let (M, g) be a compact oriented Riemannian manifold with metric tensor g of class \mathcal{C}^∞ . Then $u \in W^{1,p}(M)$ if and only if $(u \circ \psi)|_U \in W^{1,p}(\Omega)$ for all local charts $\psi : U \subset M \rightarrow \Omega \subset \mathbb{R}^n$.

We now recall the classic Sobolev embedding theorems. In light of the above proposition they can be proved by localizing and using the corresponding results for the usual Sobolev spaces in \mathbb{R}^n .

Theorem 14 (Morrey embedding). *For $p > n$ we have the continuous embedding*

$$W^{1,p}(M) \hookrightarrow \mathcal{C}(M; \mathbb{R}).$$

This embedding will be used to relate the solutions to partial differential equations involving measures to Sobolev spaces. The following trace theorem is easily proved by using a local coordinate system, flattening out ∂M , noting that the regularity of the boundary ensures that we are working in a Sobolev extension domain, and then using the usual Sobolev trace theorem.

Theorem 15 (Traces in H^1). *The restriction operator $\text{Tr} : \mathcal{C}(\overline{M}; \mathbb{R}) \rightarrow \mathcal{C}(\partial M; \mathbb{R})$ extends to a continuous operator $\text{Tr} : H^1(M) \rightarrow H^{1/2}(\partial M)$ such that integration by parts holds with the weak partial derivatives.*

Finally, we recall an elliptic regularity theorem on compact Riemannian manifolds. One can easily prove the following using the ideas from Gilbarg & Trudinger [27] and working in local coordinates. This follows since the Laplace-Beltrami operator is an elliptic differential operator (in divergence form) in any local coordinates.

Theorem 16 (Elliptic Regularity). *Let $u \in W_0^{1,p}(M)$ be such that $\Delta_g u \in W^{k,p}(M)$ for some $k \in \mathbb{N}_0$. Then $u \in W^{k+2,p}(M) \cap W_0^{1,p}(M)$ and we have the bounds*

$$\|u\|_{W^{k+2,p}(M)} \leq C \left(\|u\|_{L^p(M)} + \|\Delta_g u\|_{W^{k,p}(M)} \right),$$

for some constant $C = C(k, p) > 0$.

Chapter 3

Problem Setup

3.1 Cone Singularities and the Yamabe Equation

A Riemannian metric \tilde{g} is *conformally equivalent* to g if it can be expressed via a pointwise rescaling $\tilde{g} = e^{2u}g$ for some function $u : M \rightarrow \mathbb{R}$ called the *log conformal factor*; any such transformation preserves angles between tangent vectors, but not necessarily their length. A conformal *flattening* is any conformal rescaling such that the new metric \tilde{g} has zero Gaussian curvature. The change in curvature under a conformal rescaling is described by the *Yamabe equation*:

$$\Delta u = K - e^{2u}\tilde{K}.$$

Along the boundary ∂M , the change in geodesic curvature κ is described by the *Cherrier boundary conditions*:

$$\frac{\partial u}{\partial \mathbf{n}} = \kappa - e^u\tilde{\kappa}.$$

A derivation of these equations can be found in Aubin [6].

A conformal *cone metric* with *cone singularities* $\{p_i\}_{i=1}^m$ of *orders* $\{\alpha_i\}_{i=1}^m$ is a Riemannian metric g such that in local complex coordinates $g = e^{2u}|dz|^2$ the function

$$u(z) - \alpha_i \log |z - z(p_i)|$$

is continuous near each point p_i [59, Definition 1]. A cone metric is *polyhedral* if its Gaussian curvature is zero away from cone points (the chief example of course being the metric of any Euclidean polyhedron, as depicted in Fig. 1.1).

Now how does the Yamabe equation change in the presence of cone singularities? To understand this we will use the following generalized Gauss-Bonnet theorem.

Proposition 17 ([60, Proposition 1]). *Let (M, g) be a compact Riemann surface with conical singularities at points $\{p_j\}_{j=1}^m \subseteq M$ with order $\{\alpha_j\}_{j=1}^m$. Then*

$$\int_M K \, dA + \int_{\partial M} \kappa \, d\ell = 2\pi \left(\chi(M) + \sum_{j=1}^m \alpha_j \right),$$

where K is the Gaussian curvature of g , κ is the geodesic curvature of g , $\chi(M)$ is the Euler characteristic of M , and dA and $d\ell$ are the Riemannian volume forms on M and ∂M , respectively.

Proposition 18 ([60, Lemma 3]). *If dA_g has a conical singularity of order β at $p \in M \setminus \partial M$, then*

$$|z - a| \frac{\partial u}{\partial z} \quad \text{and} \quad |z - a| \frac{\partial u}{\partial \bar{z}} \rightarrow 0$$

when $z \rightarrow a$ where z is a local complex coordinate system in a neighborhood of p , $a = z(p)$, and $dA_g = e^{2u}|z - a|^{2\beta}|dz|^2$.

We obtain the following result, which is central to framing the cone singularity placement problem as a PDE constrained optimization.

Proposition 19. *Let (M, g) be a compact Riemann surface. Let $\tilde{g} = e^{2u}g$ be a conformal rescaling of the metric with conical singularities at points $\{p_j\}_{j=1}^m$ with orders $\{\alpha_j\}_{j=1}^m$ and zero Gaussian curvature (i.e., $\tilde{K} = 0$). Then*

$$\begin{cases} \Delta u = K + 2\pi \sum_{j=1}^m \alpha_j \delta_{p_j}, \\ \frac{\partial u}{\partial \mathbf{n}} = \kappa - e^u \tilde{\kappa}, \end{cases} \quad (3.1)$$

where $\tilde{\kappa}$ is the geodesic curvature induced by \tilde{g} .

Proof. Note that the Riemannian volume forms induced by \tilde{g} are

$$\tilde{dA} = e^{2u} dA, \quad \tilde{d\ell} = e^u d\ell,$$

where \tilde{dA} and $\tilde{d\ell}$ denote the volume forms on (M, \tilde{g}) .

From the usual Gauss-Bonnet theorem we have

$$\int_M K dA + \int_{\partial M} \kappa d\ell = 2\pi\chi(M).$$

Similarly, from Troyanov's conical Gauss-Bonnet formula (Proposition 19) we have

$$\int_M \tilde{K} \tilde{dA} + \int_{\partial M} \tilde{\kappa} \tilde{d\ell} = 2\pi\chi(M) + 2\pi \sum_{j=1}^m \alpha_j.$$

From [60, Lemma 3] we obtain that $u \in W^{1,p}(M)$ for some $p \geq 1$. Hence, by working in local coordinates and using the Sobolev trace theorem we conclude that

$$\int_M \Delta u dA = \int_{\partial M} \frac{\partial u}{\partial \mathbf{n}} d\ell.$$

We deduce from the Cherrier boundary conditions above that

$$\int_{\partial M} \tilde{\kappa} \tilde{d\ell} = \int_M \tilde{K} e^{2u} dA + \int_{\partial M} \tilde{\kappa} e^u d\ell = \int_{\partial M} \left(\kappa - \frac{\partial u}{\partial \mathbf{n}} \right) d\ell = \int_{\partial M} \kappa d\ell + \int_M \Delta u dA.$$

By subtracting the above equality from the conical Gauss-Bonnet formula we obtain

$$0 = 2\pi \sum_{j=1}^m \alpha_j + \left(2\pi\chi(M) - \int_M \kappa d\ell \right) - \int_M \Delta u dA = 2\pi \sum_{j=1}^m \alpha_j + \int_M (K - \Delta u) dA,$$

where we used the usual Gauss-Bonnet formula in the second equality. Now we see that

$$\int_M (\Delta u - K) dA = 2\pi \sum_{j=1}^m \alpha_j.$$

Now since this same result holds where M is replaced by any open subset that contains all of the cone points, and since it is zero otherwise we conclude that

$$\Delta u = K + 2\pi \sum_{j=1}^m \alpha_j \delta_{p_j}$$


in the appropriate weak sense. □

The cone angles $\theta_i := -2\pi\alpha_i$ describe the angle of opening when the flattened surface is cut through cone points p_i and laid out isometrically in the plane. We can rewrite the above equation in terms of the singular Yamabe equation in terms of the cone angles directly as

$$\Delta u = K - \sum_{j=1}^m \theta_j \delta_{p_j},$$

and this is the form of the equation we will primarily in this thesis.

Remark. To understand the relationship between the cone angle in the plane and the order of the cone singularity, we consider the prototypical case of a Euclidean cone. Consider two rays from the origin meeting with angle β (corresponding to θ_j above) in \mathbb{R}^2 (see inset figure), and let C denote the quotient metric space obtained by identifying these two rays.



We want to endow C with a Riemannian metric which induces the distance metric on C . Assume that there exists a Riemannian metric of the form

$$g = \varphi^2(r)(dr^2 + r^2d\theta^2)$$

in polar coordinates—although φ could also depend on θ , we assume it does not due to symmetry. Let $\gamma : \mathbb{S}^1 \rightarrow C$ denote the (parallel and closed) curve $r = r_0$. By measuring the length of γ with respect to this metric we have

$$L(\gamma) = \int_0^{2\pi} r_0 \varphi(r_0) d\theta = 2\pi r_0 \varphi(r_0).$$

On the other hand, since this is a curve in the plane we can measure the length of γ as the angle times the radius:

$$L(\gamma) = (2\pi - \beta) \int_0^{r_0} \varphi(r) dr.$$

Equating these and letting $M(r) = \int_0^r \varphi(s) ds$ we see

$$(2\pi - \beta)M(r) = 2\pi r M'(r),$$

and by solving this ordinary differential equation we see that $M(r) = r^{1 - \frac{\beta}{2\pi}}$. Thus,

$$\varphi(r) = M'(r) = \left(1 - \frac{\beta}{2\pi}\right) r^{-\frac{\beta}{2\pi}}.$$

Set $\beta = -2\pi\alpha$ to conclude that the Riemannian cone metric of this Euclidean cone is

$$g = (\alpha + 1)^2 r^{2\alpha} (dr^2 + r^2 d\theta^2).$$

Note that this prototypical model of a conical singularity justifies the definition of a conformal cone singularity provide earlier. Furthermore, it shows that the angle (of opening) is related to the order of the cone singularity by the relationship $\beta = -2\pi\alpha$.

3.2 Basic Problem

We recall the formulation of the optimal cones problem. Let (M, g) be a smooth surface. We want to find a pointwise conformally equivalent metric $\tilde{g} = e^{2u}g$ such that the new metric is flat away from

finite number of cone singularities. We want to choose the placement and angles of the cone singularities in an *optimal* way – that is in a way that minimizes the area distortion of the induced conformal map.

Using the singular Yamabe equation (Eqn. 3.1) we can formulate our main problem as follows:

$$\begin{array}{ll}
 \text{minimize}_{p_i \in M, \alpha_i \in \mathbb{R}} & \mathcal{E}(u) \\
 \text{subject to} & \Delta_g u = K - \sum \alpha_i \delta_{p_i} \quad \text{in } M \\
 & u = 0 \quad \text{on } \partial M,
 \end{array} \tag{P_{\text{cones}}}$$

where \mathcal{E} is a measure of area distortion induced by the conformal rescaling of the metric (to be defined in Sec. 3.3), K is the Gaussian curvature of (M, g) , and Δ_g is the Laplace-Beltrami operator on (M, g) .

3.2.1 Local Picture

A simpler, local perspective will prove very useful for both intuition and analysis—here we will assume that M is homeomorphic to the unit disk. Instead of considering the above optimization problem over the original manifold, we first consider an initial conformal flattening given by rescaling the metric by e^{2u_0} where $\Delta_g u_0 = K$ in M with $u_0|_{\partial M} \equiv 0$. (P_{cones}) then becomes

$$\begin{array}{ll}
 \text{minimize}_{p_i \in M, \alpha_i \in \mathbb{R}} & \mathcal{E}(u + u_0) \\
 \text{subject to} & \Delta_{\mathbb{R}^2} u = \sum \alpha_i \delta_{p_i} \quad \text{in } M \\
 & u = 0 \quad \text{on } \partial M,
 \end{array} \tag{3.2}$$

Note that the energy is now in terms of $u + u_0$, since if we conformally rescale the metric again, we obtain a new metric

$$\tilde{g} = e^{2u} e^{2u_0} g = e^{2(u+u_0)} g.$$

We use $\Delta_{\mathbb{R}^2}$ to denote the Laplace-Beltrami operator on $(M, e^{2u_0} g)$ to emphasize the fact that in local coordinates it is just the usual Laplacian in \mathbb{R}^2 .

From this flat point of view we get a different perspective on the problem: given the function u_0 describing the scale distortion in the initial flattening, find the best approximation by a finite weighted sum of *harmonic Green's functions*, i.e., solutions to $\Delta_{\mathbb{R}^2} u = \delta_p$ (as pictured in Fig. 3.1). This question of best approximation has a fundamentally different flavor than simply picking extrema of Gaussian curvature or peaks in the scale factor itself: it reminds us that the long tails of our harmonic Green's functions will also have an influence on the result. Hence, placing cones at local extrema may adversely affect the solution elsewhere (as sometimes occurs with greedy strategies); conversely, cones carefully arranged in flat regions can conspire to reduce distortion in regions of greater curvature (as we sometimes observe in optimal solutions). From a formal analytical viewpoint, this perspective also makes it easier to analyze the singular Yamabe equation since we can use classical results on the harmonic Green's functions for smooth subsets of \mathbb{R}^2 . Of course, for this question to be meaningful we must first answer the question of what it means for an approximation to be “best.”

3.3 Measures of Area Distortion

How do we measure the area distortion of a conformal flattening? Springborn et al. [54] remark that since a uniform global scaling changes a flattening only superficially, one can measure area distortion via the (scale-invariant) Dirichlet energy of the conformal log-scale factors:

$$\mathcal{E}_{\text{Dirichlet}}(u) = \frac{1}{2} \int_M g(\nabla u, \nabla u) dA.$$

When this energy is small, it means that scaling is near-constant, *i.e.*, low area distortion up to uniform scaling. Unfortunately, this energy is not meaningful in the context of cone flattening. To see why, consider an initial nonsingular conformal flattening via log scale factors u_0 (as in Sec. 3.2). If the subsequent scaling u satisfies

$$\Delta_{\mathbb{R}^2} u = \delta_p$$

for some point $p \in M$, then in local coordinates it is a smooth perturbation of a Euclidean harmonic Green's function, *i.e.*,

$$u(q) = \frac{1}{2\pi} \log \frac{1}{\|q - p\|} + \eta(q),$$

for some $\eta \in \mathcal{C}^\infty(M; \mathbb{R})$. One can easily show that for any ball $B(p, \varepsilon) \subset M$,

$$\int_{B(p, \varepsilon)} \|\nabla u\|^2 dA \sim \int_0^\varepsilon \frac{1}{r} dr + c(\eta) = +\infty,$$

and hence $\|\nabla(u + u_0)\|_g^2$ is not integrable. As a result, Dirichlet energy cannot distinguish between distinct configurations of cone singularities: they all have infinite energy. In the discrete case, this means that Dirichlet energy will be dependent on the resolution of mesh, since the scale factor due to cones will be better resolved—and hence much larger—by a fine mesh than a coarse one (see Fig. 3.1); a simple calculation shows that this blowup occurs at a rate of $\log(1/h)$, where h denotes the mean edge length. Any algorithm which aims to place cone singularities in a way that minimizes $\mathcal{E}_{\text{Dirichlet}}$ will therefore avoid placing singularities in densely sampled regions of a mesh—this turns out to be a practical issue for **CETM**, which drives cone angles to zero under refinement (Fig. 3.2). Moreover, since $\mathcal{E}_{\text{Dirichlet}}$ penalizes only the local change in u , it can admit large variations in scale over a domain with large diameter.

We will instead use the L^2 -norm of the log conformal factors to measure area distortion:

$$\mathcal{E}_{L^2}(u) = \int_M |u|^2 dA.$$

This energy is finite for any solution of the singular Yamabe equation, and hence converges to a finite value under mesh refinement. As noted by Myles & Zorin [45], it is also a second-order approximation of the nonlinear elastic energy [15]

$$\mathcal{E}_{\text{elastic}}(f) = \frac{1}{2} \int_M \min_{R \in SO(2)} \|df - R\|^2 dA,$$

where $f : M \rightarrow \mathbb{R}^2$ is the induced conformal flattening, and df denotes the differential of f . In computer graphics this energy is known as the *as rigid as possible* energy [53]; the energy \mathcal{E}_{L^2} is known in mechanics as the *true strain* or *Hencky strain* [30].

Unlike $\mathcal{E}_{\text{Dirichlet}}$, the L^2 energy is not invariant to constant shifts in scale. More generally, one might




Figure 3.1: Near a cone singularity, the scale factor u looks like a harmonic Green's function (*far right*), which gets significantly smoothed out when sampled onto a discrete mesh (*top row*). Methods that pick cones based on, *e.g.*, peaks of the scale factor can therefore have very different behavior on coarse and fine meshes.




Figure 3.2: Since the Dirichlet energy of a harmonic Green's function blows up under mesh refinement, cone angles obtained in CETM by setting $u = 0$ at the point of maximum distortion will tend to zero for fine meshes.

want to quotient $L^2(M)$ out by constant functions and minimize the oscillation of u , which can be written as

$$\mathcal{E}_{\text{osc}}(u) = \min_{c \in \mathbb{R}} \int_M |u - c|^2 dA = \int_M |u - \text{avg}_M(u)|^2 dA,$$

where

$$\text{avg}_M u := \frac{1}{\text{Area}(M)} \int_M u dA.$$

This energy is invariant under constant shifts in scale. This behavior is desirable since ideally we do not want to consider constant scaling as area distortion, *i.e.*, we would like $u = c$ for *any* fixed constant $c \in \mathbb{R}$. By jointly optimizing over the boundary conditions and the cone singularity placements, we find that the minimizer with \mathcal{E}_{osc} will also minimize \mathcal{E}_{L^2} .

For some applications one might also be interested in minimizing the *worst* area distortion. Unfortunately, asking to minimize area distortion in the L^∞ sense is again not meaningful in the presence of cones, since scaling goes to infinity at every cone. An interesting question for future work is to consider L^p norms for p much greater than 2, which might exhibit the desired behavior.

3.4 Relaxation

As noted in the introduction, without imposing a condition on the number of cone singularities, we can make the area distortion arbitrarily close to zero. A short proof of this fact using the theory of optimal transportation is given in Corollary 44. If we fix the number of cone singularities that we want to place then the optimization problem (P_{cones}) becomes akin to combinatorial optimization problem, and we do not expect to find any efficient algorithms for solving it directly. This motivates us to look for a relaxed optimization problem that properly discretizes the underlying smooth optimization problem and is efficiently solvable.

Note that we can rewrite (P_{cones}) as

$$\begin{aligned}
& \underset{p_i \in M, \alpha_i \in \mathbb{R}}{\text{minimize}} && \mathcal{E}(u) \\
& \text{subject to} && \mu = \sum \alpha_i \delta_{p_i} \\
& && \Delta_g u = K - \mu \quad \text{in } M \\
& && u = 0 \quad \text{on } \partial M,
\end{aligned} \tag{3.3}$$

The idea is to drop the hard constraint $\mu = \sum \alpha_i \delta_{p_i}$, and instead incorporate it into the energy.

The naïve way to incorporate this soft-constraint into the energy is to use the sparsity inducing L^1 -regularization. The resulting optimization problem reads

$$\begin{aligned}
& \underset{f \in L^1(M)}{\text{minimize}} && \mathcal{E}(u) + \lambda \|f\|_{L^1} \\
& \text{subject to} && \Delta_g u = K - f \quad \text{in } M \\
& && u = 0 \quad \text{on } \partial M,
\end{aligned} \tag{P_{L^1-relaxed}}$$

where $\lambda > 0$ is a specified tuning parameter. However, there are both serious practical and theoretical issues with using the L^1 -norm to recover cone singularity configurations. In the smooth setting the existence of minimizers to $(P_{L^1-relaxed})$ does not hold in general. More precisely, for a solution to exist there must be a minimizing sequence with a convergent subsequence, *i.e.*, a sequence that is at least weakly precompact in $L^1(M)$. The Dunford-Pettis theorem implies that a sequence is weakly precompact if and only if it is equi-integrable [26, Theorem 2.54]. However, neither a generic minimizing sequence nor a sequence converging in L^1 to a Dirac delta measure will be equi-integrable. Hence, even though the discrete, finite dimensional ℓ^1 problem has solutions, they are not concentrated at isolated vertices, nor stable with respect to tessellation (see Fig. 3.3). Furthermore, since this problem arises in the smooth setting a different choice of finite elements does not fix this problem.

As discussed in Sec. 2.2, the natural generalization of the L^1 -space is the space of measures. Moreover, this space exhibits the necessary compactness properties to guarantee the existence of minimizers! We find that the sparsity-inducing property of the L^1 -norm naturally carries over to the $\mathcal{M}(Y)$ -norm, and the numerical issues discussed above do not arise when the optimization problem over the space of measures is treated appropriately. To summarize our approach: rather than optimizing over the placement and angles of the cone singularities, we optimize over all finite signed Radon measures μ and add a sparsity promoting regularization. We introduce a tuning parameter $\lambda \geq 0$, and relax (P_{cones}) to obtain

$ \begin{aligned} & \underset{\mu \in \mathcal{M}(M)}{\text{minimize}} && \mathcal{E}(u) + \lambda \ \mu\ _{\mathcal{M}(M)} \\ & \text{subject to} && \Delta_g u = K - \mu \quad \text{in } M \\ & && u = 0 \quad \text{on } \partial M, \end{aligned} $	(P_{relaxed})
--	------------------------

Using the dual characterization of $\mathcal{M}(Y)$ presented in Sec. 2.2 we will be able to employ Fenchel-Rockafellar duality to obtain equivalent optimization problems framed over function spaces. These reformulations will lend themselves to efficient numerical techniques for determining the optimal placement of cone singularities. A simple and illustrative comparison between the results obtained from the L^1 regularized $(P_{L^1-relaxed})$ and the measure norm regularized (P_{relaxed}) is presented in Fig. 3.3.

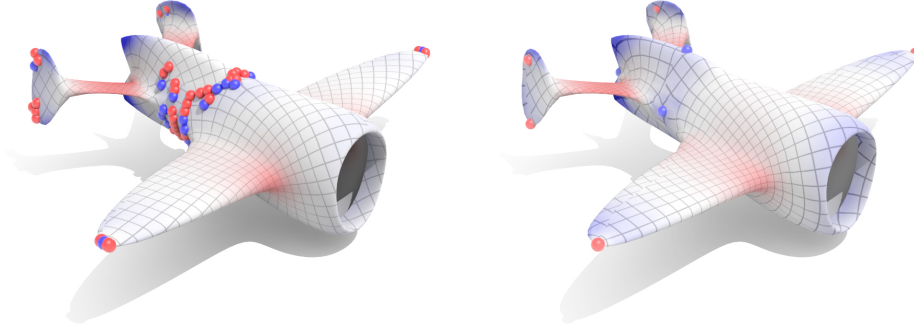


Figure 3.3: *Left:* naively discretizing measures as piecewise linear functions and using L^1 regularization yields solutions that have low area distortion but are not concentrated at isolated points. *Right:* By considering a dual problem formulated in terms of ordinary functions, we obtain a much better numerical scheme for computing isolated cones.

3.5 Generalizations

In this section, we generalize the basic formulation to allow greater control over how and where cone singularities are distributed. For this purpose we consider two distinct regions $U_{\mathcal{E}}$, an open subset of M , and $U_{\mathcal{R}}$, a compact subset of \overline{M} . We will only measure the area distortion over $U_{\mathcal{E}}$, only place cone singularities in $U_{\mathcal{R}}$. This is useful for placing cones only in, say, visible regions (Sec. 6.2.1), or for only optimizing for prescribed boundary conditions (Sec. 6.2.3); see examples in Figures 6.8 and 3.4.

Furthermore, we introduce two positive weight functions $w_{\mathcal{E}}, w_{\mathcal{R}} \in \mathcal{C}^\infty(M; \mathbb{R}_{>0})$ that are bounded away from zero. Now we consider the generalized energy

$$\mathcal{E}_{L^2}^{w_{\mathcal{E}}}(u) := \frac{1}{2} \int_{U_{\mathcal{E}}} |u|^2 w_{\mathcal{E}} dA,$$

and the generalized regularization

$$\mathcal{R}^{w_{\mathcal{R}}}(\mu) := \int_{U_{\mathcal{R}}} w_{\mathcal{R}} d|\mu|.$$

Here $|\mu|$ is the total variation measure associated to μ . By re-weighting the energy and the measure space regularization we incentivize our optimization problem to place cone singularities in the regions where $w_{\mathcal{R}}$ is small in a way that minimizes the area distortion in the regions where $w_{\mathcal{E}}$ is large. As described in Sec. 6.2.1, this reweighting can be used to promote the placement of cone singularities in geometrically meaningful locations; a natural reweighting of the area distortion is given by a local feature size, which we express as a function of the Gaussian curvature. We will often denote the terms $\mathcal{E}_{L^2}^{w_{\mathcal{E}}}$ and $\mathcal{R}^{w_{\mathcal{R}}}$ as \mathcal{E} and \mathcal{R} , respectively.

We also extend the basic formulation to allow us to optimize over the boundary conditions as well. One can either optimize over the Dirichlet boundary conditions, or alternatively one can optimize over the geodesic curvature on the boundary by using the Cherrier formula. By using the Cherrier boundary conditions we obtain a clear geometric interpretation of adding the measure-regularization on the boundary. This generalization is treated simply, both analytically and algorithmically, by encoding the boundary conditions in the solution operator (see Sec. 4.2 and Sec. 5.1).

Finally, we generalize our formulation to allow us to optimize over singularity configurations with only positive cone angles. To achieve this we let $\mathcal{M}^+(U_{\mathcal{R}})$ denote the set of positive Radon measures.




Figure 3.4: We can selectively restrict cone placement to any user-specified region. Here for instance, by shooting rays (*top left*) we can determine the region visible from a particular point of view (*bottom left*). If we now restrict our search to this region—while still penalizing distortion in the front—singularities that would ordinarily appear on the front (*center*) instead get “pushed” to the back (*right*).

Now we set ι_+ the associated indicator functions of this set. That is,

$$\iota_+(\mu) := \begin{cases} 0 & \text{if } \mu \in \mathcal{M}^+(U_\mu), \\ +\infty & \text{else.} \end{cases}$$

Remark. We can also optimize over the space of all negative measures, but in practice we find that many models prefer no cone singularities over any configurations with only negative cone angles.

Our relaxed optimization problem now reads

$$\begin{array}{l} \underset{\substack{\mu \in \mathcal{M}(U_\mathcal{R}) \\ h \in H^1(M)}}{\text{minimize}} \quad \frac{1}{2} \int_{U_\mathcal{E}} |u|^2 w_\mathcal{E} \, dA + \lambda \int_{U_\mathcal{R}} w_\mathcal{R} \, d|\mu| \\ \text{subject to} \quad \Delta_g u = K - \mu \quad \text{in } M \\ \quad \quad \quad u = h \quad \quad \quad \text{on } \partial M. \end{array} \quad (P_{\text{generalized}})$$

For simplicity and clarity of presentation we restrict much of our analysis to zero Dirichlet boundary conditions. However, in our construction of the solution operator we will use Robin (mixed) boundary conditions since well-posedness of the PDE with measure-valued right hand sides is extremely simple. Of course, the underlying optimization problem is the same since we can consider the above optimization (with either choice of boundary control) over all $u \in W^{1,p}(M)$ for suitably chosen p .

Chapter 4

Theoretical Analysis

In this chapter we will analyze the optimization problem from Chapter 3 in the continuous setting, formulate an appropriate optimization problem through convex duality, and analyze properties of the minimizing measures. The setup and basic analysis of the elliptic control problem formulated over the space of measures essentially follows the ideas presented by Clason & Kunisch [17, 18]. Our treatment of the problem varies primarily in how we formulate the abstract framework and in the properties of the minimizers that we care about and study. The main motivation of the abstract formulation is that a number of the optimization problems [14, 17, 18] formulated over non-reflexive Banach spaces can be treated in a single unified framework described below.

We organize the chapter as follows: we begin by presenting a fairly abstract optimization problem, before specializing the framework to our basic optimization problem. Formulating the abstract control problem distills and emphasizes the main features and assumptions of the cone singularity placement problem. Next, we discuss the solution operator for the cones problem—that is, we study solutions to Poisson equation with measure-valued right hand sides. In Sec. 4.3 we specialize the abstract framework from Sec. 4.1 to our relaxed problem, and derive an optimality system in two ways: (1) through Fenchel-Rockafellar duality, and (2) directly through Lagrange duality. These two approaches are related, and help emphasize different aspects of our relaxed control problem. In Sec. 4.4 we prove some properties of the minimizing measures. In particular, we show that they *cannot* represent any cone singularities; however, we show that they can be rounded to cones in a way that does not significantly change the area distortion. Finally, we introduce and analyze a sequence of regularized problems which will be useful for the purposes of discretization.

4.1 Abstract Framework

For the abstract framework, we consider the constrained optimization problem

$$\begin{array}{ll} \text{minimize} & \mathcal{E}(u) + \mathcal{R}(\mu) \\ \mu \in \mathcal{M}, u \in U & \\ \text{subject to} & s(u, \mu) = 0 \text{ in } X^*. \end{array} \quad (P_{\text{abstract}})$$

Here, U and \mathcal{M} are Banach spaces. The state and the control variables are related through the map $s : U \times \mathcal{M} \rightarrow X^*$, where X is yet another Banach space.

For (P_{abstract}) to be well-posed, we assume the following:

Assumption 1. $\mathcal{M} \cong \mathcal{C}^*$, where \mathcal{C} is a separable Banach space.

Assumption 2. X and U are reflexive.

Assumption 3. $\mathcal{E} : U \rightarrow \mathbb{R}$ is bounded from below and \mathcal{C}^1 Fréchet differentiable.

Assumption 4. $\mathcal{R} : \mathcal{M} \rightarrow \overline{\mathbb{R}}$ is convex, proper, and weak- * sequentially lower semicontinuous. Furthermore, \mathcal{R} is bounded from below and coercive.

Assumption 5. For every $\mu \in \mathcal{M}$ there exists a unique $u \in U$ such that

$$s(u, \mu) = 0 \quad \text{in } X^*.$$

The associated solution operator $S : \mathcal{M} \rightarrow U$ given by $S\mu := u$ is affine, bounded, and weak- * to strongly sequentially continuous, *i.e.*,

$$S\mu_n \rightarrow S\mu \quad \text{in } U \quad \text{for all } \mu_n \xrightarrow{*} \mu \quad \text{in } \mathcal{M}.$$

Using the solution operator, we can define reduced objectives for (P_{abstract}) . We define the reduced energy as $e : \mathcal{M} \rightarrow \mathbb{R}$ as

$$e(\mu) := \mathcal{E}(S\mu),$$

and the reduced objective $v : \mathcal{M} \rightarrow \overline{\mathbb{R}}$ as

$$v(\mu) := e(\mu) + \mathcal{R}(\mu) = \mathcal{E}(S\mu) + \mathcal{R}(\mu).$$

Note that e is sequentially weak- * to strongly continuous since the energy is continuous and since the solution operator is continuous.

We endow \mathcal{M} with the weak- * topology, $\sigma(\mathcal{M}, \mathcal{C})$. Since \mathcal{C} is separable, the unit ball in \mathcal{M} with the weak- * topology is metrizable, and thus sequentially precompact. In light of our assumptions and this compactness property, a standard application of the direct method of the calculus of variations gives us existence of minimizers.

Theorem 20. *The constrained optimization problem (P_{abstract}) admits at least one globally optimal solution $(\bar{\mu}, \bar{u}) = (\bar{\mu}, S\bar{\mu}) \in \mathcal{M} \times U$.*

Proof. First, note that

$$\bar{v} := \inf_{\mu \in \mathcal{M}} v(\mu) \in \mathbb{R}$$

since v is bounded from below and since \mathcal{R} is proper.

Now consider any minimizing sequence $\{\mu_n\}_{n=1}^{\infty} \subseteq \mathcal{M}$. For n large enough we have that

$$\mathcal{R}(\mu_n) \leq v(\mu_n) < \bar{v} + 1.$$

Since \mathcal{R} is coercive there exists $c \in \mathbb{R}$ satisfying

$$\|\mu_n\|_{\mathcal{M}} \leq c$$

for all $n \in \mathbb{N}$. By the Banach-Alaoglu theorem we obtain a subsequence (not relabeled) such that $\mu_n \xrightarrow{*} \bar{\mu}$ for some $\bar{\mu} \in \mathcal{M}$. Since v is weak- * sequentially lower semicontinuous we conclude that

$$v(\bar{\mu}) \leq \liminf_{n \rightarrow \infty} v(\mu_n) = \bar{v},$$

which implies that $\bar{\mu} \in \mathcal{M}$ is the minimizer since we also have that $\bar{v} \leq v(\bar{\mu})$. \square

To determine the first order optimality conditions in this general framework, we will be unable to use a black box like Fenchel-Rockafellar convex duality. However, we will still be able to characterize the minimizers directly using the subdifferential. Since \mathcal{M} may *not* be reflexive, it is natural to define the subdifferential with respect to the weak- * topology. That is, the subdifferential (denoted by $\underline{\partial}$ to

differentiate it from the usual subdifferential) of $\mathcal{R} : \mathcal{M} \rightarrow \mathbb{R}$ is a subset of \mathcal{C} (instead of \mathcal{M}^*):

$$\underline{\partial}\mathcal{R}(\mu) := \{\psi \in \mathcal{C} : \langle \tilde{\mu} - \mu, \psi \rangle_{\mathcal{M} \times \mathcal{C}} \leq \mathcal{R}(\mu) \leq \mathcal{R}(\tilde{\mu}) \text{ for all } \tilde{\mu} \in \mathcal{M}\}.$$

Note that e is Fréchet differentiable since S and \mathcal{E} are as well. For the sake of generality (this may not hold in general!), we assume that the differential, de , is represented by an element in \mathcal{C} in the sense that

$$de(\bar{\mu})[v] = \langle v, \nabla e(\bar{\mu}) \rangle_{\mathcal{M} \times \mathcal{C}}$$

for some $\nabla e(\bar{\mu}) \in \mathcal{C}$. As is usual in the calculus of variations and Riemannian geometry, we call this the *gradient* of e at $\bar{\mu}$. In many cases this gradient will be characterized as the solution of an appropriate adjoint PDE.

With the above setup, we go on to deduce the necessary first order optimality system for (P_{abstract}) .

Proposition 21. *Let $\bar{\mu} \in \mathcal{M}$ be a minimizer of (P_{abstract}) . Then $-\nabla e(\bar{\mu}) \in \underline{\partial}\mathcal{R}(\bar{\mu})$.*

Proof. Since $\bar{\mu} \in \mathcal{M}$ is a minimizer we see that

$$v(\bar{\mu}) \leq v(\bar{\mu} + \varepsilon(\mu - \bar{\mu}))$$

for all $\mu \in \mathcal{M}, \varepsilon > 0$. Furthermore, since \mathcal{R} is convex we deduce

$$0 \leq v(\bar{\mu} + \varepsilon(\mu - \bar{\mu})) - v(\bar{\mu}) \leq e(\bar{\mu} + \varepsilon(\mu - \bar{\mu})) - e(\bar{\mu}) + \varepsilon(\mathcal{R}(\bar{\mu}) - \mathcal{R}(\mu)).$$

Now by dividing by ε and taking $\varepsilon \rightarrow 0^+$ we have

$$0 \leq \lim_{\varepsilon \rightarrow 0^+} \frac{e(\bar{\mu} + \varepsilon(\mu - \bar{\mu})) - e(\bar{\mu})}{\varepsilon} + (\mathcal{R}(\bar{\mu}) - \mathcal{R}(\mu)) = de(\bar{\mu})[\mu - \bar{\mu}] + (\mathcal{R}(\bar{\mu}) - \mathcal{R}(\mu)).$$

Rearranging and using the representation of $\nabla e(\bar{\mu})$ in \mathcal{C} we have

$$\langle \mu - \bar{\mu}, -\nabla e(\bar{\mu}) \rangle + \mathcal{R}(\mu) \leq \mathcal{R}(\bar{\mu}),$$

which is exactly $-\nabla e(\bar{\mu}) \in \underline{\partial}\mathcal{R}(\bar{\mu})$. □

This simple abstract reformulation turns out to be incredibly expressive for a large class of PDE constrained optimizations. In particular, it can be used to properly formulate both parabolic and elliptic optimal control problems that appear in the literature.

4.2 Poisson Equation with Measures

Recall that the singular Yamabe equation is simply a Poisson equation with a measure valued right hand side. Since our relaxation is formulated over the space of Radon measures, we need to study the existence and regularity of solutions to the partial differential equation

$$\Delta u = \mu, \tag{4.1}$$

for $\mu \in \mathcal{M}(M)$ and appropriate boundary conditions.

Weak Solutions For the weak formulation of Eqn. 4.1 with zero Dirichlet boundary conditions we consider the bilinear form

$$a(u, \varphi) = (\nabla u, \nabla \varphi)_{L^2(M; TM)} = \int_M g(\nabla u, \nabla \varphi) dA$$

for $u, \varphi \in H_0^1(M)$. Note that a can be continuously extended to a bilinear form $a : W_0^{1,p}(M) \times W_0^{1,p'}(M) \rightarrow \mathbb{R}$ by Hölder's inequality for any $p \in (1, +\infty)$ where $\frac{1}{p} + \frac{1}{p'} = 1$. Recall that by the Sobolev embedding theorem we have for $p' > 2$ that $W^{1,p'}(M)$ continuously embeds into the space of Hölder continuous functions. Thus,

$$W_0^{1,p'}(M) \hookrightarrow \mathcal{C}_0(M)$$

for $p' > 2$. In the following we fix $p' > 2$.

Definition 22. For $\mu \in \mathcal{M}(M)$ we say that $u \in W_0^{1,p}(M)$ is a *weak solution* to Eqn. 4.1 if

$$a(u, \varphi) = \langle \mu, \varphi \rangle_{\mathcal{M} \times \mathcal{C}} = \int_M \varphi \, d\mu$$

for all $\varphi \in W_0^{1,p'}(M)$.

In dimensions $n \geq 2$ the classical variational framework is not equipt to deal with measure-valued right hand sides. This essentially comes down to dimensionality considerations from the Sobolev embedding theorems. We present two distinct approaches to obtain solutions to the above PDE—the first is given by Green's potentials, and is useful for proving a regularity result on the minimizing measure. Alternatively, we will construct a solution operator via duality, which will be useful for *directly* applying convex duality (this is not necessary if we are to immediately apply the abstract framework).

4.2.1 Solutions by Green's Potentials

When obtaining solutions to the Poisson equation through Green's potentials it is most natural to consider *Dirichlet* boundary conditions $b \in H^1(M)$.

Remark. We just need b to be the trace of some function $H^1(M)$. I know for regions in \mathbb{R}^N that this space is simply the fractional Sobolev space $H^{1/2}(\partial M)$. However, I am not familiar with the definitions of the fractional Sobolev spaces on abstract Riemannian manifolds. Therefore, I just consider boundary conditions $b \in H^1(M)$, which automatically encodes this information.

Dealing with the boundary conditions $b \in H^1(M)$ poses no real challenge. Let $u_b \in H^1(M)$ solve

$$\begin{cases} \Delta u_b = 0 & \text{in } M, \\ u_b = b & \text{on } \partial M, \end{cases}$$

where the boundary conditions are understood in the trace sense. The existence of u_b follows from classical variational arguments:

Proposition 23. *There exists a unique $u_b \in H^1(M)$ solving the above Dirichlet boundary problem.*

Proof. Let $f := -\Delta b \in H^{-1}(M)$. Consider the bilinear form $a : H_0^1(M) \times H_0^1(M) \rightarrow \mathbb{R}$ given by

$$a(u, \varphi) = \int_M g(\nabla u, \nabla \varphi) \, dA,$$

which is clearly both bounded (continuous) and coercive. Now by the Lax-Milgram theorem we obtain a unique $\tilde{u}_b \in H^1(M)$ such that

$$a(\tilde{u}_b, \varphi) = \langle f, \varphi \rangle_{H^{-1} \times H^1}$$

for all $\varphi \in H_0^1(M)$. Now let $u_b := \tilde{u}_b + b$ to see that

$$a(u_b, \varphi) = \int_M g(\nabla \tilde{u}_b, \nabla \varphi) dA + \int_M g(\nabla b, \nabla \varphi) dA = \langle f, b \rangle_{H^{-1} \times H^1} + \int_M g(\nabla b, \nabla \varphi) dA = 0.$$

□

It is now straightforward to obtain a solution to $\Delta u = \mu$ with $u|_{\partial M} = b$: we simply need to solve

$$\begin{cases} \Delta v = \mu & \text{in } M, \\ v = 0 & \text{on } \partial M, \end{cases}$$

and then set $u := v + u_b$. By the linearity of the trace operator, and by the linearity of the Laplace-Beltrami operator, we see that u solves Eqn. 4.1 with $u|_{\partial M} = b$.

To obtain v solving the PDE from the previous page, we use the Green's function on M . In particular, let G_M denote the Green's function of the Laplace-Beltrami operator on (M, g) :

Theorem 24 (Aubin [6, Theorem 4.17]). *Let \bar{M} be an oriented compact Riemannian manifold with boundary of class \mathcal{C}^∞ . There exists, G_M , the Green's function of the Laplace-Beltrami operator, which has the following properties*

(a) All functions $\varphi \in \mathcal{C}^2(\bar{M})$ satisfy

$$\varphi(p) = \int_M G_M(p, q) \Delta \varphi(q) dA(q) - \int_{\partial M} \nu^i \nabla_{iq} G_M(p, q) \varphi(q) d\ell(q),$$

where ν is the unit normal vector oriented to the outside and $d\ell$ is the Riemannian volume form on ∂M corresponding to the Riemannian metric $\iota^* g$ where $\iota : \partial M \rightarrow \bar{M}$ is the canonical embedding.

(b) G_M is \mathcal{C}^∞ on $\bar{M} \times \bar{M}$ minus the diagonal.

(c) For $p, q \in \bar{M}$ and $C = C(\text{dist}(p, \partial M))$ the following hold

$$\begin{aligned} |G_M(p, q)| &< C(1 + |\log r|), \\ |\nabla_q G_M(p, q)| &< \frac{C}{r}, \\ |\nabla_q^2 G_M(p, q)| &< \frac{C}{r^2}. \end{aligned}$$

(d) $G(p, q) > 0$ for $p, q \in M$.

(e) $G(p, q) = G(q, p)$.

From (a) of the above theorem we see that $\mu \in \mathcal{M}(M)$ that

$$v(\mathbf{x}) := \int_M G_M(\mathbf{x}, \mathbf{y}) d\mu(\mathbf{y})$$

satisfies $\Delta v = \mu$ with $v = 0$ on ∂M in the sense of distributions (and even in the weak sense, although it is a bit more work to show that $v \in W_0^{1,p}(M)$). The regularity of distributional solutions to the Poisson equation follows from the elliptic regularity of the Laplace-Beltrami operator on (M, g) .

4.2.2 Solutions Through Duality

One can alternatively construct solutions to the Poisson equation with measure-valued right hand sides through duality. This method can be traced back to Stampacchia [55]. The construction we present

here can be found in Clason & Schiela [19], although the notations and setup we use are restricted to the precise elliptic control problem from our work.

The main idea is to construct an operator $\Delta_\star : \mathcal{C}(M) \rightarrow W_0^{1,p}(M)^\star$ for some $p < 2$ and then obtaining an isomorphism $\Delta := (\Delta_\star)^\star$ as the Banach space adjoint of Δ_\star . In this section we will primarily consider zero Dirichlet boundary conditions, since adding boundary conditions in $H^1(M)$ follows in the same way as before. The main benefit of constructing the solution operator through duality is that it becomes clear how to use this operator for purposes of convex duality.

If one is to incorporate boundary conditions into this solution operator it becomes most natural to consider the *Robin boundary conditions* since one immediately obtains well-posedness of solutions with measures supported on the boundary—it is not clear how to show such well-posedness for Neumann boundary conditions, and in light of the Sobolev trace theorem it does not even make sense for Dirichlet boundary conditions.

Solution Operator Let $\text{dom}(\Delta_g)_\star$ denote the maximal subset of $H_0^1(M)$ where the operator

$$\begin{aligned} (\Delta_g)_\star : \text{dom}(\Delta_g)_\star &\rightarrow W_0^{1,p'}(M)^\star =: W^{-1,p'}(M) \\ \varphi &\mapsto (u \mapsto ((\Delta_g)_\star \varphi)(u) := a(u, \varphi)). \end{aligned}$$

is an isomorphism. It follows from the classical L^p -elliptic regularity theory of the Laplace-Beltrami operator that $\text{dom}(\Delta_g)_\star = W^{1,p}(M)$. Since $W^{1,p}(M)$ is dense in $\mathcal{C}(\overline{M})$, we can consider the Banach space adjoint $\Delta_g := ((\Delta_g)_\star)^\star$.

$$\Delta_g : \text{dom} \Delta_g \subset W^{1,p'}(M) \rightarrow \mathcal{M}(\overline{M}),$$

where the maximal domain of definition $\text{dom} \Delta_g$ is given by

$$\text{dom} \Delta_g := \{u \in W^{1,p'}(M) : \exists c_u \in \mathbb{R} \text{ with } ((\Delta_g)_\star \varphi)(u) \leq c_u \|\varphi\|_{\mathcal{C}(\overline{M})}, \forall \varphi \in \text{dom}(\Delta_g)_\star\}.$$

To understand this operator directly note that for any $u \in \text{dom} \Delta_g$, the linear map $\varphi \mapsto a(u, \varphi)$ defines the element of the dual space $\mathcal{C}(\overline{M})^\star$:

$$\Delta_g u = a(u, \cdot).$$

By the Riesz representation theorem, $\Delta_g u$ can be identified with an element of $\mathcal{M}(\overline{M})$. By [12, Theorem 2.21], the weak Laplace-Beltrami operator Δ_g is closed and invertible since the pre-adjoint $(\Delta_g)_\star$ is.

Remark (Robin Boundary Conditions). Fix $\mu \in \mathcal{M}(\overline{M})$ and consider the PDE

$$\begin{cases} \Delta u = \mu|_M & \text{in } M, \\ \frac{\partial u}{\partial \mathbf{n}} + u = \mu|_{\partial M} & \text{on } \partial M. \end{cases}$$

Instead of considering the usual bilinear form a from before, we can consider the continuous and coercive bilinear form $\tilde{a} : H^1(M) \times H^1(M) \rightarrow \mathbb{R}$ by

$$\tilde{a}(u, \varphi) := (\nabla u, \nabla \varphi)_{L^2(M; TM)} + (u, \varphi)_{L^2(\partial M)} = \int_M g(\nabla u, \nabla \varphi) dA + \int_{\partial M} u \varphi d\ell.$$

To see that \tilde{a} encodes a weak formulation of the above PDE note that for $u, \varphi \in \mathcal{C}^\infty(M)$, by integration by parts,

$$\tilde{a}(u, \varphi) = \int_M (\Delta u) \varphi dA + \int_{\partial M} \left(\frac{\partial u}{\partial \mathbf{n}} + u \right) \varphi d\ell.$$

Now if $\tilde{a}(u, \varphi) = \langle \mu, \varphi \rangle$ for some $\mu \in \mathcal{C}^\infty(M)$ for all suitable test functions φ then we see that the above PDE is satisfied in the classical sense. As before, we continuously extend \tilde{a} to a bilinear form $W^{1,p}(M) \times W^{1,p'}(M) \rightarrow \mathbb{R}$ and define the solution operator above in the same way. Note that the integration by parts does not change sign since by convention the Laplace-Beltrami operator is chosen to be positive-semidefinite (unlike the usual Laplace operator in \mathbb{R}^n).

To deal with the generalizations where we restrict the domains where we measure area distortion and place cone singularities we introduce the following operator.

Extension by Zero Operator We now introduce the restriction operator

$$E_\star : \mathcal{C}(\overline{M}) \rightarrow \mathcal{C}(U_{\mathcal{R}}), \quad (E_\star \varphi)(\mathbf{x}) = \varphi(\mathbf{x}) \text{ for all } \mathbf{x} \in U_{\mathcal{R}}.$$

Again, we utilize the Riesz representation theorem to consider the adjoint $E := (E_\star)^\star$ as a mapping

$$E : \mathcal{M}(U_{\mathcal{R}}) \rightarrow \mathcal{M}(\overline{M}).$$

It is straightforward to see that E is the canonical extension by zero operator: for $\mu \in \mathcal{M}(U_{\mathcal{R}})$ and $\varphi \in \mathcal{C}(M)$ we have

$$\int_{U_{\mathcal{R}}} E_\star \varphi \, d\mu = \int_M \varphi \, d(E\mu).$$

Generalized Solution Operator Finally, we define the weak- \star dual of the generalized solution operator (with either zero Dirichlet or Robin boundary conditions)

$$S_\star : W^{1,p'}(M)^\star \rightarrow \mathcal{C}(U_{\mathcal{R}}), \quad h \mapsto E_\star(\Delta_g)_\star^{-1}h.$$

As before, set $S = (S_\star)^\star$:

$$S : \mathcal{M}(U_{\mathcal{R}}) \rightarrow W^{1,p'}(M).$$

Since all of these operators are bounded, it trivially follows that $S = \Delta_g^{-1}E$, and S is closed, i.e. if $\mu_n \xrightarrow{\star} \mu$ in $\mathcal{M}(U_{\mathcal{R}})$ and $S\mu_n \rightarrow u$ in $W^{1,p'}(M)$ for some $h \in W^{1,p'}(M)$ then $S\mu = u$.

4.3 Existence and Duality for Optimal Cones

Recall the relaxed cone singularity placement problem

$$\begin{array}{ll} \text{minimize}_{\mu \in \mathcal{M}(U_{\mathcal{R}})} & \mathcal{E}(u) + \mathcal{R}(\mu) \\ \text{subject to} & \Delta_g u = \Omega - \mu \quad \text{in } M \end{array} \quad (P_{\text{relaxed}})$$

where $\Omega := K \, dA$ is the curvature 2-form on (M, g) . Here \mathcal{E} and \mathcal{R} are the generalized notions of area distortion and regularization from Sec. 3.5.

We begin by showing that this problem satisfies all of the assumptions of the abstract framework (Sec. 4.1) to deduce existence of solutions. For this purpose we set

$$\mathcal{M} := \mathcal{M}(U_{\mathcal{R}}), \quad U := W^{1,p}(M), \quad \mathcal{C} := \mathcal{C}(U_{\mathcal{R}}), \quad \text{and} \quad X := W^{1,p'}(M).$$

We encode the constraint $\Delta u = \Omega - \mu$ through the state equation $s : W^{1,p}(M) \times \mathcal{M}(U_{\mathcal{R}}) \rightarrow W^{1,p'}(U_{\mathcal{R}})^\star$

defined by

$$\langle s(u, \mu), \varphi \rangle_{W^{1,p'}(M)^* \times W^{1,p'}(M)} := a(u, \varphi) - \langle (\Omega - \mu) \lfloor_{U_{\mathcal{R}}}, \varphi \rangle,$$

where $a : W^{1,p}(M) \times W^{1,p'}(M) \rightarrow \mathbb{R}$ is the bilinear form associated to the Laplace-Beltrami operator.

Proposition 25. *The regularization $\mathcal{R} : \mathcal{M}(U_{\mathcal{R}}) \rightarrow \mathbb{R}$ is convex, proper, bounded from below, coercive, and weak- * sequentially lower semicontinuous.*

Proof. Convexity of \mathcal{R} follows from the convexity of the total variation norm (i.e., $|\theta\mu + (1-\theta)\nu|(E) \leq \theta|\mu|(E) + (1-\theta)|\nu|(E)$ for all Borel subsets E). Since $\mathcal{R}(\mathbf{0}) = 0$ we see that it is proper. Since $\mathcal{R}(\mu) \geq 0$ we see that it is bounded from below. For any $\mu \in \mathcal{M}(U_{\mathcal{R}})$ we have that

$$\mathcal{R}(\mu) \geq \left(\inf_{x \in U_{\mathcal{R}}} w_{\mathcal{R}}(x) \right) \|\mu\|_{\mathcal{M}(U_{\mathcal{R}})},$$

and so we see that \mathcal{R} is coercive. Since $\mathcal{M}(U_{\mathcal{R}}) \cong \mathcal{C}(U_{\mathcal{R}})^*$ by the Banach-Alaoglu theorem (weak- * compactness of the unit ball) we conclude that the norm $\|\cdot\|_{\mathcal{M}(U_{\mathcal{R}})}$ is weak- * lower-semicontinuous. To see that the reweighted regularization is also lower-semicontinuous consider a sequence $\mu_n \xrightarrow{*} \mu$ and define

$$\nu_n(E) := \int_E w_{\mathcal{R}} d\mu_n, \quad \nu(E) := \int_E w_{\mathcal{R}} d\mu.$$

Note that

$$\langle \nu_n - \nu, \varphi \rangle_{\mathcal{M} \times \mathcal{C}} = \int_{U_{\mathcal{R}}} \varphi d(\nu_n - \nu) = \int_{U_{\mathcal{R}}} \varphi \cdot w_{\mathcal{R}} d(\mu_n - \mu) \rightarrow 0$$

as $n \rightarrow +\infty$ for all $\varphi \in \mathcal{C}(U_{\mathcal{R}})$. Hence, $\nu_n \xrightarrow{*} \nu$. Since $\|\nu_n\|_{\mathcal{M}(U_{\mathcal{R}})} = \mathcal{R}(\mu_n)$ we conclude that \mathcal{R} is also weak- * sequentially lower semicontinuous since the norm is.

To restrict the optimization to the space of positive measures we need to prove that indicator function ι_+ is weak- * sequentially lower semicontinuous. This follows easily since

$$\iota_+(\mu) \neq +\infty \iff \langle \mu, \varphi \rangle_{\mathcal{M} \times \mathcal{C}} \geq 0$$

for all continuous $\varphi \in \mathcal{C}(U_{\mathcal{R}})$ satisfying $\varphi \geq 0$ pointwise. \square

Proposition 26. *The area distortion energy $\mathcal{E} : W^{1,p}(M) \rightarrow \mathbb{R}$ is continuously Fréchet differentiable and bounded from below.*

Proof. Let $u, v \in W^{1,p}(M)$. Since $W^{1,p}(M) \hookrightarrow L^2(M, dA)$ we can use the Lebesgue dominated convergence theorem to conclude that

$$\frac{\partial \mathcal{E}}{\partial v}(u) = \lim_{t \rightarrow 0} \frac{1}{2} \int_{U_{\mathcal{E}}} \left(\frac{|u + tv|^2 - |u|^2}{t} \right) w_{\mathcal{E}} dA = \int_{U_{\mathcal{E}}} uvw_{\mathcal{E}} dA.$$

Since the mapping $v \mapsto \frac{\partial \mathcal{E}}{\partial v}(u)$ is clearly a continuous linear functional on $W^{1,p}(M)$ we see that \mathcal{E} is Gâteaux differentiable. Furthermore, the above expression for the Gâteaux differential shows that the mapping

$$\begin{aligned} d\mathcal{E} : W^{1,p}(M) &\rightarrow W^{1,p}(M)^* \\ \tilde{u} &\mapsto d\mathcal{E}(\tilde{u}) \end{aligned}$$

is continuous at u . Hence, \mathcal{E} is Fréchet differentiable at u with the Fréchet differential coinciding with the Gâteaux differential. Furthermore, \mathcal{E} is continuously Fréchet differentiable. Since $\mathcal{E}(u) \geq 0$ we also have that \mathcal{E} is bounded from below. \square

Remark. In fact, \mathcal{E} is continuously Fréchet differentiable in the weaker space $L^2(M, dA)$.

Theorem 27 (Existence of minimizers). *The relaxed optimization (P_{relaxed}) admits a global minimizer.*

Proof.

- Assumption 1 holds since $\mathcal{M}(M) \cong \mathcal{C}(M)^*$.
- Assumption 2 holds by basic properties (i.e., Riesz representation theorem) of Sobolev spaces.
- Assumption 3 holds by Proposition 26.
- Assumption 4 holds by Proposition 25.
- Assumption 5 holds by the choice of the state equation $s : W^{1,p}(M) \times \mathcal{M}(M) \rightarrow W^{1,p'}(M)$ and the construction of the solution operator from Sec. 4.2.

Now the existence of a global minimizer is a direct corollary of Theorem 20. \square

4.3.1 Applying the Abstract Framework

We now use the abstract framework to derive the first order optimality system described as variational inequalities directly through Lagrange duality. Furthermore, by applying the abstract framework from Sec. 4.1 we do not need to make any change of variables to make the solution operator linear. Instead, we consider the solution operator $S_\Omega(\mu) := S(\Omega - \mu)$, where again $\Omega = K dA$ is the curvature 2-form (note that S_Ω is the solution operator needed in Assumption 5).

Consider the Lagrangian $\mathcal{L} : \mathcal{M}(U_{\mathcal{R}}) \times W^{1,p}(M) \times W^{1,p'}(M) \rightarrow \mathbb{R}$ defined as

$$\begin{aligned} \mathcal{L}(\mu, u, \varphi) &:= \mathcal{E}(u) - \langle s(u, \mu), \varphi \rangle_{W^{1,p'}(M)^* \times W^{1,p'}(M)} \\ &= \mathcal{E}(u) - a(u, \varphi) + \langle (\Omega - \mu) \lrcorner_{U_{\mathcal{R}}}, \varphi \rangle_{\mathcal{M} \times \mathcal{C}} \\ &= \int_{U_{\mathcal{E}}} |u|^2 dA - \int_M g(\nabla u, \nabla \varphi) dA + \int_{U_{\mathcal{R}}} \varphi \cdot K dA - \int_{U_{\mathcal{R}}} \varphi d\mu. \end{aligned}$$

Now by construction of the solution operator S_Ω we see that

$$e(\mu) = \mathcal{L}(\mu, S_\Omega \mu, \varphi)$$

for any $\mu \in \mathcal{M}(U_{\mathcal{R}})$ and $\varphi \in W^{1,p'}(M)$.

Now fix any measure $\hat{\mu} \in \mathcal{M}(U_{\mathcal{R}})$ and consider the corresponding scale factors $\hat{u} = S_\Omega \hat{\mu}$. Using the chain rule to differentiate the reduced objective e we get

$$de(\hat{\mu})[v] = d\mathcal{L}_u(\hat{\mu}, \hat{u}, \varphi)[w] + d\mathcal{L}_\mu(\hat{\mu}, \hat{u}, \varphi)[v]$$

for any variation $v \in \mathcal{M}(U_{\mathcal{R}})$, where $w = S_\Omega v$. To simplify the above we take $\hat{\varphi} \in W^{1,p'}(M)$ to be a weak solution of the adjoint equation

$$d\mathcal{L}_u(\hat{\mu}, \hat{u}, \hat{\varphi})[\psi] = d\mathcal{E}(\hat{u})[\psi] - a(\psi, \hat{\varphi}) = 0$$

for all $\psi \in W^{1,p}(M)$. Thus,

$$de(\hat{\mu})[v] = d\mathcal{L}_\mu(\hat{\mu}, \hat{u}, \hat{\varphi})[v] = -\langle v \lrcorner_{U_{\mathcal{R}}}, \hat{\varphi} \rangle_{\mathcal{M} \times \mathcal{C}}$$

for all $v \in \mathcal{M}(U_{\mathcal{R}})$. We can express these computations concisely as

Proposition 28. *The gradient of the reduced energy $e = \mathcal{E} \circ S_\Omega$ is given by*

$$\nabla e(\hat{\mu}) = -\hat{\varphi},$$

where $\varphi \in \mathcal{C}(U_{\mathcal{R}})$ solves the adjoint equation

$$a(\psi, \hat{\varphi}) = d\mathcal{E}(\hat{u})[\psi]$$

for all $\psi \in W^{1,p}(M)$.

We now express the adjoint equation as a weak formulation of a partial differential equation:

Proposition 29. *The solution to the adjoint equation is a weak solution to $\Delta_g \hat{\varphi} = w_\varepsilon \hat{u}$.*

Proof. Expanding out $a(\psi, \hat{\varphi}) = d\mathcal{E}(\hat{u})[\psi]$ and using the Riesz identification of $L^2(M)$ with its dual we see that

$$\int_M g(\nabla \hat{\varphi}, \nabla \psi) dA = \int_M \hat{u} \psi w_\varepsilon dA$$

Hence, $\Delta_g \hat{\varphi} = w_\varepsilon \hat{u}$ as desired. \square

Summarizing, if $(\bar{\mu}, \bar{u}) = (\bar{\mu}, S(\Omega - \bar{\mu}))$ is the optimal solution of our relaxed optimization problem. Then there exists a unique adjoint state $\bar{\varphi}$ solving the adjoint state equation

$$a(\psi, \bar{\varphi}) = d\mathcal{E}(\bar{u})[\psi]$$

for all $\psi \in W^{1,p}(M)$. Expanding the adjoint equation in terms of differential operators, this means that there exists a unique adjoint state $\bar{\varphi} \in W^{1,p'}(M)$ such that

$$\Delta \bar{\varphi} = w_\varepsilon \bar{u},$$

in the weak sense. Furthermore, by Proposition 21, $\bar{\varphi}$ satisfies the variational inequality

$$\langle \mu - \bar{\mu}, \bar{\varphi} \rangle + \mathcal{R}(\bar{\mu}) \leq \mathcal{R}(\mu)$$

for all $\mu \in \mathcal{M}(U_{\mathcal{R}})$. This is equivalent to $\bar{\varphi} \in \partial \mathcal{R}(\bar{\mu})$.

We can summarize the entire first order optimality system as follows:

Proposition 30. *Let $\bar{\mu} \in \mathcal{M}(U_{\mathcal{R}})$ denote the minimizing measure. Let \bar{u} and $\bar{\varphi}$ denote the minimizing states. Then the following holds.*

$$\begin{cases} \Delta \bar{u} = \Omega - \bar{\mu}, \\ \Delta \bar{\varphi} = w_\varepsilon \bar{u}, \\ \bar{\varphi} \in \partial \mathcal{R}(\bar{\mu}) \end{cases} \quad (OS_{\text{cones}})$$

where the above Poisson equations are understood in the appropriately weak sense.

4.3.2 Applying Fenchel-Rockafellar duality

Since our relaxed optimization problem is *convex* we can directly apply Fenchel-Rockafellar duality to obtain an equivalent optimization problem over the usual function spaces. However, since $\mathcal{M}(M)$ is not reflexive, we will need to formulate a *pre-dual* problem that achieves our relaxed optimization problem as it's Fenchel-Rockafellar dual. We will use the following result to compute the Fenchel conjugates.

Proposition 31 ([47, Theorem 23.5]). *Let V be a Banach space, and let $\mathcal{F} : V \rightarrow \overline{\mathbb{R}}$ be a proper convex function, $v \in V$, and $L \in V^*$. Then $\mathcal{F}^*(L) = \langle L, v \rangle_{V^*, V} - \mathcal{F}(v) \iff L \in \partial \mathcal{F}(v)$.*

To apply Fenchel-Rockafellar duality directly (cf. Theorem 2), we need to slightly reformulate (P_{relaxed}) . By making the change of variables $\sigma := \Omega - \mu$, the solution operator $\sigma \mapsto S\sigma = u$ becomes linear (instead of just affine linear). With this change, the optimization above now reads:

$$\begin{array}{ll} \text{minimize}_{\sigma \in \mathcal{M}(U_{\mathcal{R}})} & \mathcal{E}(u) + \mathcal{R}(\Omega - \sigma) \\ \text{subject to} & \Delta u = \sigma \end{array} \quad (\tilde{P}_{\text{relaxed}})$$

Theorem 32. *The relaxed optimization problem $(\tilde{P}_{\text{relaxed}})$ is the dual of*

$$\begin{array}{ll} \text{minimize}_{u \in L^2(M, dA)} & \frac{1}{2} \int_{U_{\mathcal{E}}} \frac{|u|^2}{w_{\mathcal{E}}} dA - \int_{U_{\mathcal{R}}} \varphi \cdot \Omega \\ \text{subject to} & \Delta_* \varphi = u \\ & |\varphi(\mathbf{x})| \leq \lambda w_{\mathcal{R}}(\mathbf{x}) \text{ for all } \mathbf{x} \in U_{\mathcal{R}}. \end{array} \quad (4.2)$$

and the minimizers are related by the following optimality system

$$\begin{cases} \Delta \bar{u} = \Omega - \bar{\mu}, \\ \Delta \bar{\varphi} = w_{\mathcal{E}} \bar{u}, \\ \bar{\mu} \llcorner_{U_{\mathcal{R}}} \in \partial i_{\lambda w_{\mathcal{R}}}(\bar{\varphi}). \end{cases} \quad (OS_{\text{cones}})$$

Proof. Define $\mathcal{F} : L^2(U_{\mathcal{E}}, dA) \rightarrow \overline{\mathbb{R}}$ by

$$\mathcal{F}(u) := \mathcal{E}\left(\frac{u}{w_{\mathcal{E}}}\right) = \frac{1}{2} \int_{U_{\mathcal{E}}} \frac{|u|^2}{w_{\mathcal{E}}} dA.$$

From Proposition 26 and the chain rule we have that $d\mathcal{F}(u) = u/w_{\mathcal{E}}$, where, as usual, we use the Riesz identification of $L^2(M, dA) \cong L^2(M, dA)^*$. Now using Proposition 31 we deduce that $\mathcal{F}^* : L^2(U_{\mathcal{E}}, dA) \rightarrow \overline{\mathbb{R}}$ is given by

$$\mathcal{F}^*(u) = \frac{1}{2} \int_{U_{\mathcal{E}}} |u|^2 w_{\mathcal{E}} dA = \mathcal{E}(u).$$

Similarly, define $\mathcal{G} : \mathcal{C}(U_{\mathcal{R}}) \rightarrow \overline{\mathbb{R}}$ by

$$\mathcal{G}(\varphi) := - \int_{U_{\mathcal{R}}} \varphi \cdot \Omega + \begin{cases} 0 & \text{if } |\varphi(\mathbf{x})| \leq \lambda w_{\mathcal{R}}(\mathbf{x}) \text{ for all } \mathbf{x} \in U_{\mathcal{R}}, \\ +\infty & \text{else.} \end{cases}$$

Using the Riesz representation theorem for $\mathcal{C}(U_{\mathcal{R}})$ we consider $\mathcal{G}^* : \mathcal{M}(U_{\mathcal{R}}) \rightarrow \overline{\mathbb{R}}$, and directly compute for $\mu \in \mathcal{M}(U_{\mathcal{R}})$

$$\mathcal{G}^*(\mu) = \sup \left\{ \int_{U_{\mathcal{R}}} \varphi d(\mu + \Omega) : \varphi \in \mathcal{C}(U_{\mathcal{R}}), |\varphi(\mathbf{x})| \leq \lambda w_{\mathcal{R}}(\mathbf{x}) \right\}.$$

Let $U_{\mathcal{R}}^+$ denote the support of $((\mu + \Omega) \llcorner_{U_{\mathcal{R}}})^+$ and $U_{\mathcal{R}}^-$ denote the support of $((\mu + \Omega) \llcorner_{U_{\mathcal{R}}})^-$. Recall that by the Jordan decomposition theorem these sets are disjoint. Now construct a sequence of continuous

functions $\{\varphi_n\}_{n=1}^\infty \subset \mathcal{C}(U_{\mathcal{R}})$ such that pointwise

$$\begin{cases} \varphi_n(\mathbf{x}) \rightarrow \lambda w_{\mathcal{R}}(\mathbf{x}) & \text{for all } \mathbf{x} \in U_{\mathcal{R}}^+, \\ \varphi_n(\mathbf{x}) \rightarrow -\lambda w_{\mathcal{R}}(\mathbf{x}) & \text{for all } \mathbf{x} \in U_{\mathcal{R}}^-, \end{cases}$$

as $n \rightarrow \infty$. By Fatou's lemma we have that

$$\mathcal{G}^*(\mu) \geq \liminf_{n \rightarrow \infty} \int_{U_{\mathcal{R}}} \varphi_n d(\mu + \Omega) = \lambda \int_{U_{\mathcal{R}}} w_{\mathcal{R}} d|\mu + \Omega|$$

The reverse inequality is immediate from the inequality constraints on φ . Summarizing, we have shown that \mathcal{G}^* is the measure space regularization appearing in our relaxed formulation.

We use the adjoint of the injection $W^{1,p}(M) \hookrightarrow L^2(M)$ to identify $L^2(M)$ with a subset of $\text{dom } S_\star = W^{1,p}(M)^\star$. Now a direct application of Fenchel-Rockafellar duality (Theorem 2) shows that

$$\inf_{v \in L^2(M, dA)} \mathcal{F}(v) + \mathcal{G}(S_\star v) = - \min_{\sigma \in \mathcal{M}(U_{\mathcal{R}})} \mathcal{F}^*(S\sigma) + \mathcal{G}^*(-\sigma).$$

By the computations above

$$\mathcal{F}^*(S\sigma) + \mathcal{G}^*(-\sigma) = \mathcal{E}(u) + \lambda \mathcal{R}(\Omega - \sigma),$$

where $\Delta u = \sigma$. This shows the equivalence of our two formulations. Furthermore, the optimality conditions from the theorem also follow from the Fenchel-Rockafellar duality theorem. \square

Remark. To enforce the condition that the measures are strictly positive or negative we simply replace the indicator function in \mathcal{G} with

$$\begin{cases} 0 & \text{if } \pm \varphi(\mathbf{x}) \leq \lambda w_{\mathcal{R}}(\mathbf{x}) \text{ for all } \mathbf{x} \in U_{\mathcal{R}}, \\ +\infty & \text{else.} \end{cases}$$

An identical computation reveals that

$$\mathcal{G}_\pm^*(\mu) = \lambda \int_{U_{\mathcal{R}}} w_{\mathcal{R}} d|\mu + \Omega| + \iota_\pm(\mu),$$

as desired. Here \mathcal{G}_\pm denotes \mathcal{G} with the corresponding indicator function above.

We claim that the variational inequality from the abstract formulation coincides with the subdifferential inclusion from the optimality system obtained from Theorem 32.

Proposition 33. Let $\mu \in \mathcal{M}(U_{\mathcal{R}})$ and $\varphi \in \mathcal{C}(U_{\mathcal{R}})$. Then

$$\varphi \in \partial \mathcal{R}(\bar{\mu}) \iff \mu \in \partial i_{w_{\mathcal{R}}}(\varphi).$$

Proof. By using Proposition 31 we see

$$\mu \in \partial i_{w_{\mathcal{R}}}(\varphi) \iff i_{w_{\mathcal{R}}}(\varphi) + i_{w_{\mathcal{R}}}^*(\mu) = \langle \mu, \varphi \rangle \iff \mathcal{R}(\mu) = \langle \mu, \varphi \rangle$$

since $i_{w_{\mathcal{R}}}(\varphi) = 0$ and $i_{w_{\mathcal{R}}}^* = \mathcal{R}$. For any $\tilde{\mu} \in \mathcal{M}(U_{\mathcal{R}})$ we compute

$$\langle \tilde{\mu}, \varphi \rangle = \int_{U_{\mathcal{R}}} \varphi(\mathbf{x}) d\tilde{\mu}(\mathbf{x}) \leq \int_{U_{\mathcal{R}}^+} w_{\mathcal{R}}(\mathbf{x}) d\tilde{\mu}(\mathbf{x}) + \int_{U_{\mathcal{R}}^-} w_{\mathcal{R}}(\mathbf{x}) d\tilde{\mu}(\mathbf{x}) = \mathcal{R}(\tilde{\mu}),$$

where $U_{\mathcal{R}}^\pm = U_{\mathcal{R}} \cap \text{supp}(\tilde{\mu}^\pm)$. Adding the $\mathcal{R}(\mu) = \langle \mu, \varphi \rangle$ to the above gives us

$$\langle \tilde{\mu} - \mu, \varphi \rangle + \mathcal{R}(\mu) \leq \mathcal{R}(\tilde{\mu})$$

for all $\tilde{\mu} \in \mathcal{M}(U_{\mathcal{R}})$. Thus, by definition, $\bar{\varphi} \in \partial \mathcal{R}(\mu)$ as desired. \square

So we see that the direct approach using the Lagrangian gives the same optimality system as the application of Fenchel-Rockafellar duality.

The main motivation for using the direct approach is directly applicable to a much larger class of problems (*i.e.*, S is nonlinear or \mathcal{E} is not convex). On the other hand, the approach through convex duality provides a simple black-box approach and emphasizes the intimate connection between sparse optimization and *state constrained optimization*.

Another fundamental difference between the two approaches is that the Fenchel-Rockafellar duality approach required us to formulate an optimization problem over the *pre-dual* space, whereas the direct approach was easily formulated over the non-reflexive Banach space $\mathcal{M}(M)$ by simply considering the subdifferential with respect to the appropriate (weak-*) topology.

Characterizing the Minimizers We now want to exactly characterize the support of the minimizers using only the first order optimality system.

Proposition 34. *Let $\bar{\mu}$ be the optimizing measure, and let $\bar{\varphi}$ be the optimal adjoint state. We have that*

$$\begin{aligned} \text{supp } \bar{\mu}^+ &\subseteq \{\mathbf{x} \in U_{\mathcal{R}} : \bar{\varphi}(\mathbf{x}) = +w_{\mathcal{R}}(\mathbf{x})\}, \\ \text{supp } \bar{\mu}^- &\subseteq \{\mathbf{x} \in U_{\mathcal{R}} : \bar{\varphi}(\mathbf{x}) = -w_{\mathcal{R}}(\mathbf{x})\}. \end{aligned}$$

Proof. This result follows from the optimality conditions presented in (OS_{cones}) . In particular, we will use the condition that $\bar{\mu} \in \partial i_{\lambda w_{\mathcal{R}}}(\bar{\varphi} \chi_{U_{\mathcal{R}}})$. From Proposition 31 and this subdifferential relationship we have that

$$\langle \bar{\mu}, \bar{\varphi} \chi_{U_{\mathcal{R}}} \rangle = i_{\lambda w_{\mathcal{R}}}(\bar{\varphi} \chi_{U_{\mathcal{R}}}) + i_{\lambda w_{\mathcal{R}}}^*(\bar{\mu}) = \int_M w_{\mathcal{R}}(\mathbf{x}) d|\bar{\mu}|(\mathbf{x}). \quad (4.3)$$

In the second equality we use that fact that since $i_{\lambda w_{\mathcal{R}}}$ is subdifferentiable at $\bar{\varphi} \chi_{U_{\mathcal{R}}}$ that $\bar{\varphi} \chi_{U_{\mathcal{R}}}$ is in the effective domain of $i_{\lambda w_{\mathcal{R}}}$. Note that $\bar{\mu} \ll |\bar{\mu}|$, and so the Radon-Nikodym derivative $\text{sgn } \bar{\mu} := d\bar{\mu}/d|\bar{\mu}|$ is in $L^1(M, d|\bar{\mu}|)$. Using this and Eqn. 4.3 we obtain

$$\int_{U_{\mathcal{R}}} \text{sgn } \bar{\mu}(\mathbf{x}) \bar{\varphi} \chi_{U_{\mathcal{R}}}(\mathbf{x}) d|\bar{\mu}|(\mathbf{x}) = \int_{U_{\mathcal{R}}} w_{\mathcal{R}}(\mathbf{x}) d|\bar{\mu}|(\mathbf{x}).$$

So we conclude that $\bar{\varphi} \chi_{U_{\mathcal{R}}}(\mathbf{x}) \text{sgn } \bar{\mu}(\mathbf{x}) = \lambda w_{\mathcal{R}}(\mathbf{x})$ for $|\bar{\mu}|$ -almost-every $\mathbf{x} \in U_{\mathcal{R}}$, it now immediately follows that

$$\begin{aligned} \text{supp } \bar{\mu}^+ &\subseteq \{\mathbf{x} \in U_{\mathcal{R}} : \text{sgn } \bar{\mu} = +1\} \subseteq \{\mathbf{x} \in M : \bar{\varphi}(\mathbf{x}) = +w_{\mathcal{R}}(\mathbf{x})\}, \\ \text{supp } \bar{\mu}^- &\subseteq \{\mathbf{x} \in U_{\mathcal{R}} : \text{sgn } \bar{\mu} = -1\} \subseteq \{\mathbf{x} \in M : \bar{\varphi}(\mathbf{x}) = -w_{\mathcal{R}}(\mathbf{x})\}, \end{aligned}$$

\square

This proposition provides some intuition regarding why the minimizing measure should be sparse. If $\bar{\varphi}$ was real analytic and $w_{\mathcal{R}}$ was constant then we would immediately obtain that $\bar{\mu}$ consists of only Dirac measures. Of course, we shouldn't expect $\bar{\varphi}$ to have such high regularity; nevertheless, the main takeaway is that we expect $\bar{\varphi}$ and $w_{\mathcal{R}}$ to only coincide on lower dimensional submanifolds of M . Of course, there is no guarantee that these functions will only coincide on such lower dimensional sets.

4.4 Regularity and Rounding

In this section we aim to more thoroughly understand the sparsity structure of the minimizing measures—in particular, we want to answer the question of whether the relaxed optimization problem provides us with cone singularities. Using our construction of the solution operator through Green’s potentials we prove the negative result that the minimizing measures never represent any cone singularities. To rectify this situation, we show that “rounding” the minimizing measures to linear combinations of Dirac deltas does not substantially increase the area distortion.

In this section, we make the assumption that $U_{\mathcal{R}} = U_{\mathcal{E}} = M$. For simplicity we assume that the boundary conditions are fixed to $u|_{\partial M} = 0$ and that $w_{\mathcal{R}} \equiv 1$. We also assume that M is homeomorphic to the unit disk in \mathbb{R}^2 to allow us to perform an initial flattening—let $u_0 \in \mathcal{C}^\infty(M)$ be the log conformal factors which make $(M, e^{2u_0}g)$ flat. Using this initial conformal flattening, we can consider M as a subset of \mathbb{R}^2 as described in Sec. 3.2.1.

Using the local picture, we can consider M as an open subset of \mathbb{R}^2 and the question is to optimally approximate the scale factors $u_0 : M \rightarrow \mathbb{R}$ through a sum of harmonic Green’s functions. That is, we can write our relaxed cones optimization using the initial log-conformal factors as

$$\boxed{\begin{array}{ll} \text{minimize} & \mathcal{E}(u + u_0) + \mathcal{R}(\mu) \\ \mu \in \mathcal{M}(M) & \\ \text{subject to} & \Delta u = \mu \quad \text{in } M \end{array}} \quad (P_{\text{local}})$$

Here \mathcal{E} is reweighted by e^{2u_0} to ensure that our formulation is conformally invariant:

$$\mathcal{E}(u) := \int_M |u|^2 w_{\mathcal{E}} dA = \int_M |u|^2 w_{\mathcal{E}} e^{2u_0} dx.$$

The regularization on the other hand is measured on the original (non-flattened) manifold. Now we prove some small results regarding the structure of the minimizers.

Proposition 35. *Let $\bar{\mu} \in \mathcal{M}(M)$ be the optimizing measure. Then $\text{supp } \bar{\mu}$ is compact in M . In particular,*

$$\text{supp } \bar{\mu} \subseteq M_\delta := \{x \in M : \text{dist}(x, \partial M) > \delta\}$$

for some $\delta > 0$. Furthermore, $\text{dist}(\text{supp } \bar{\mu}^+, \text{supp } \bar{\mu}^-) > \delta$.

Proof. Recall that the adjoint state $\bar{\varphi} \in \mathcal{C}(M)$ by the Sobolev embedding $W_0^{1,p'}(M) \hookrightarrow \mathcal{C}(M)$. Since $\bar{\varphi}|_{\partial M} \equiv 0$ and since $w_{\mathcal{R}}$ is bounded away from zero, we deduce that there exists some $\delta_1 > 0$ such that

$$|\bar{\varphi}(x)| < \frac{1}{2} w_{\mathcal{R}}(x)$$

for all $x \in M \setminus M_{\delta_1}$. Now from Proposition 34 we deduce that $\text{supp } \bar{\mu}^+ \subseteq M_{\delta_1}$.

Similarly, since $w_{\mathcal{R}}$ is bounded away from zero, we can find $\eta > 0$ such that $w_{\mathcal{R}}(x) > \eta$ for all $x \in M$. Proposition 34 yields $\bar{\varphi}(x) \geq \eta$ in $\text{supp } \bar{\mu}^+$ and $\bar{\varphi}(x) \leq -\eta$ in $\text{supp } \bar{\mu}^-$. By the continuity of $\bar{\varphi}$ we deduce that there exists some $\delta_2 > 0$ such that

$$\text{dist}(\{x \in M : \bar{\varphi}(x) \geq +\eta\}, \{x \in M : \bar{\varphi}(x) \leq -\eta\}) > \delta_2.$$

Therefore, $\text{dist}(\text{supp } \bar{\mu}^+, \text{supp } \bar{\mu}^-) > \delta_2$. The desired result follows by taking $\delta = \min\{\delta_1, \delta_2\}$. \square

The following proposition guarantees that the maximal distortion obtained by solving our optimization problem will be less than the distortion from the initial conformal flattening—this will be used to prove that $\bar{u} \in H^1(M)$, which implies the optimizing measure cannot contain Dirac deltas.

Proposition 36. *Let $\bar{\mu}$ be the optimizing measure of (P_{local}) , and let $\Delta \bar{u} = \bar{\mu}$ with $\bar{u} = 0$ on ∂M . Then $\bar{u} \in L^\infty(M)$ with*

$$\|\bar{u}\|_{L^\infty(M)} \leq \|u_0\|_{L^\infty(M)} + \left(\inf_{\mathbf{x} \in M} w_\mathcal{E}(\mathbf{x}) e^{2u_0(\mathbf{x})} \right)^{-1}.$$

Proof. As before, consider the adjoint state $\bar{p} \in \mathcal{C}(M)$ that satisfies $\Delta_{\mathbb{R}^2} \bar{p} = \bar{u} - u_0$ with $\bar{p} = 0$ on ∂M . Assume, for the sake of contradiction, $\|\bar{u}\|_{L^\infty(M)} > \|u_0\|_{L^\infty(M)} + \left(\inf_{\mathbf{x} \in M} w_\mathcal{E}(\mathbf{x}) e^{2u_0(\mathbf{x})} \right)^{-1}$. Thus, there exists $\varepsilon > 0$ such that the set

$$A := \left\{ \mathbf{x} \in M : |\bar{u}(\mathbf{x})| > \|u_0\|_{L^\infty(M)} + \left(\inf_{\mathbf{x} \in M} w_\mathcal{E}(\mathbf{x}) e^{2u_0(\mathbf{x})} \right)^{-1} + \varepsilon \right\}$$

has positive measure. Without loss of generality, we can assume that $\sup_{\mathbf{x} \in M} \bar{u}(\mathbf{x}) > \|u_0\|_{L^\infty(M)} + \left(\inf_{\mathbf{x} \in M} w_\mathcal{E}(\mathbf{x}) e^{2u_0(\mathbf{x})} \right)^{-1} + \varepsilon$. By the maximum principle for logarithmic potentials [48, Corollary 3.3] there exists some $\mathbf{x}_0 \in \text{supp } \bar{\mu}^+$ satisfying

$$\bar{u}(\mathbf{x}_0) > \|u_0\|_{L^\infty(M)} + \left(\inf_{\mathbf{x} \in M} w_\mathcal{E}(\mathbf{x}) e^{2u_0(\mathbf{x})} \right)^{-1} + \varepsilon.$$

Now let $\delta > 0$ be such that $\text{dist}(\text{supp } \bar{\mu}^+, \text{supp } \bar{\mu}^-) > \delta$ (see Proposition 35). Note that A is open since \bar{u} is lower-semicontinuous. We deduce that $A \cap B(\mathbf{x}_0, \delta)$ is open. Thus, we can find some $0 < r < \delta$ such that $B(\mathbf{x}_0, r) \subset A \cap B(\mathbf{x}_0, \delta)$. Now since $\mathbf{x}_0 \in \text{supp } \bar{\mu}^+$ we have that $\bar{\varphi}(\mathbf{x}_0) = \lambda > 0$. Now let $y \in \mathcal{C}^\infty(M)$ be a solution to

$$\begin{cases} -\Delta y = -\varepsilon & \text{in } B(\mathbf{x}_0, r), \\ y = 0 & \text{on } \partial B(\mathbf{x}_0, r). \end{cases}$$

By the maximum principle for the Laplace operator, we deduce that $y(\mathbf{x}_0) < 0$. Now set $Y := -\bar{\varphi} - y$. Using the adjoint equation for $\bar{\varphi}$ as well as the lower bound of \bar{u} in $B(\mathbf{x}_0, r)$ we see that

$$\begin{cases} -\Delta Y = -e^{2u_0} w_\mathcal{E}(\bar{u} - u_0) + \varepsilon \leq 0 & \text{in } B(\mathbf{x}_0, r), \\ Y = \bar{\varphi} & \text{on } \partial B(\mathbf{x}_0, r). \end{cases}$$

Now by the maximum principle for subharmonic functions we conclude that $\sup_{\mathbf{x} \in B(\mathbf{x}_0, r)} Y(\mathbf{x}) = Y(\mathbf{x}_{\max}) = \bar{\varphi}(\mathbf{x}_{\max})$ for some $\mathbf{x}_{\max} \in \partial B(\mathbf{x}_0, r)$. We conclude by noting that

$$\bar{\varphi}(\mathbf{x}_{\max}) = Y(\mathbf{x}_{\max}) \geq Y(\mathbf{x}_0) = \bar{\varphi}(\mathbf{x}_0) - y(\mathbf{x}_0) > \bar{\varphi}(\mathbf{x}_0) = \lambda.$$

This is a contradiction with the adjoint state constraint $|\bar{\varphi}(\mathbf{x})| \leq \lambda$ for all $\mathbf{x} \in M$. So we conclude that $\|\bar{u}\|_{L^\infty(M)} \leq \|u_0\|_{L^\infty(M)} + \left(\inf_{\mathbf{x} \in M} w_\mathcal{E}(\mathbf{x}) e^{2u_0(\mathbf{x})} \right)^{-1}$, as desired. \square

Remark. *It is unclear whether the above result holds when considering non-constant weighting $w_\mathcal{R}$. The above argument is based on arriving to a contradiction with the adjoint state constraints $|\varphi(\mathbf{x})| \leq \lambda w_\mathcal{R}(\mathbf{x})$. Thus, the challenge with having a reweighted regularization is in making sure that $w_\mathcal{R}(\mathbf{x}_0) > w_\mathcal{R}(\mathbf{x}_{\max})$. Alternatively, we can try to reweight the functions y and Y to incorporate $\sup w_\mathcal{R}(\mathbf{x})$, but I could not see how to make this work consistently.*

To see why the L^∞ bound will show that $\bar{u} \in H_0^1(M)$ we first consider the case when $\bar{\mu} \in \mathcal{C}^\infty(M)$.

Lemma 37. *Let $f \in \mathcal{C}_c^\infty(M)$, and consider the measure μ given by $d\mu = f \, dA$. Then*

$$\int_M u(\mathbf{x}) \, d\mu(\mathbf{x}) = \int_M \|\nabla u(\mathbf{x})\|_g^2 \, dA(\mathbf{x}),$$

where $\Delta u = \mu$ with $u = 0$ on ∂M .

Proof. Since f is smooth, it is clear that u is smooth. Integration by parts reveals

$$\int_M u(\mathbf{x}) d\mu(\mathbf{x}) = \int_M u(\mathbf{x})f(\mathbf{x}) dA(\mathbf{x}) = - \int_M u(\mathbf{x})\Delta u(\mathbf{x}) dA(\mathbf{x}) = \int_M \|\nabla u(\mathbf{x})\|_g^2 dA(\mathbf{x}).$$

Here we used the fact that u vanishes on ∂M in the integration by parts to obtain an especially simple expression for the $H^1(M)$ -norm. \square

By an approximation argument we will be able to show that the above result holds even when μ is not absolutely continuous with respect to the Riemannian volume form.

Theorem 38. *The optimizing measure $\bar{\mu} \in H_0^1(M)^*$.*

Proof. Since $\Delta \bar{u} = \bar{\mu}$ we have that

$$\int_M \bar{u}(\mathbf{x}) d\bar{\mu}(\mathbf{x}) \leq \|\bar{u}\|_{L^\infty(M)} \|\bar{\mu}\|_{\mathcal{M}}.$$

From Proposition 36 we see that $\|\bar{u}\|_{L^\infty(M)}$ is finite, and so the above expression is finite.

To show that the left hand side is the L^2 -norm of $\nabla \bar{u}$ we use Lemma 37 and a simple approximation argument: Let η_ε be the family of standard mollifiers, and define $\mu_\varepsilon := \eta_\varepsilon * \bar{\mu}$. Then we know that $\mu_\varepsilon \in \mathcal{C}^\infty(M; \mathbb{R})$ and $\mu_\varepsilon d\mathbf{x} \rightarrow \bar{\mu}$ in $\mathcal{M}(M)$ as $\varepsilon \rightarrow 0^+$. Set $u_\varepsilon := S(\mu_\varepsilon d\mathbf{x})$. Then by Lemma 37 we have that

$$\|\nabla u_\varepsilon\|_{L^2(M; TM)}^2 = \int_M u_\varepsilon d\mu_\varepsilon.$$

Now by standard properties of mollifiers we have

$$\lim_{\varepsilon \rightarrow 0^+} \|\nabla u_\varepsilon\|_{L^2(M; TM)}^2 = \|\nabla \bar{u}\|_{L^2(M; TM)}^2;$$

in particular,

$$\|\nabla \bar{u}\|_{L^2(M; TM)}^2 = \lim_{\varepsilon \rightarrow 0^+} \|\nabla u_\varepsilon\|_{L^2(M; TM)}^2 = \liminf_{\varepsilon \rightarrow 0^+} \int_M u_\varepsilon d\mu_\varepsilon \leq \int_M \bar{u} d\bar{\mu}.$$

Thus,

$$\|\nabla \bar{u}\|_{L^2(M; TM)}^2 \leq \|\bar{u}\|_{L^\infty(M)} \|\bar{\mu}\|_{\mathcal{M}} < +\infty.$$

The Poincaré inequality shows $\bar{u} \in H_0^1(M)$. Recalling the isomorphism $\Delta : H_0^1(M) \rightarrow H^{-1}(M)$ we deduce $\bar{\mu} \in H^{-1}(M)$. \square

This technical result is important since it tells us that the optimizing measure cannot consist of any Dirac deltas. This is because the Dirac delta measure is not in $H_0^1(M)^$ since $\dim M = 2$. Intuitively, this is because every Dirac delta (or cone singularity) can be rounded to a slightly “smoothed” out version in a way that brings down the area distortion in a very small way.*

Although we do not have Dirac deltas in $H_0^1(M)^*$, numerically we find that the minimizers are very close to Dirac deltas. We now want to *round* the minimizing measure to a collection of delta measures. To begin to describe a rounding procedure we need to describe what it means for two measures to be close.

Optimal transportation provides the natural framework to discuss the similarity of measures. Recall that the *Wasserstein distance* between two probability measures μ, ν is given by

$$d_{\mathcal{W}_2}(\mu, \nu) := \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{M \times M} \|\mathbf{x} - \mathbf{y}\|^2 d\pi(\mathbf{x}, \mathbf{y}) \right)^{1/2}.$$

We will show that the H^{-2} -norm is *exactly* our area distortion measure. Hence, to show that *rounding* is appropriate we need to show that the Wasserstein spaces embed continuously into $H^{-2}(M)$. Recall that $H^{-2}(M)$ is the dual space $H_0^2(M)^*$. Using Lagrange multipliers and the definition of the dual norm we obtain the following:

Proposition 39. *Let $\mu \in \mathcal{M}(M) \hookrightarrow H^{-2}(M)$. Then*

$$\|\mu\|_{H^{-2}(M)} = \left(\int_M |u|^2 d\mathbf{x} \right)^{1/2}$$

where $\Delta u = \mu$ in M with $u|_{\partial M} \equiv 0$.

We first prove our embedding result when the domain M is simply a ball in \mathbb{R}^2 . In what follows, Φ the fundamental solution of the Laplace operator in \mathbb{R}^2 .

Theorem 40. *Let $R > 0$ and let $\mathcal{P}(\overline{B(0, R)})$ denote the space of probability measures supported in $\overline{B(0, R)}$ in \mathbb{R}^2 . Then the identity map*

$$\text{id} : \left(\mathcal{P}(\overline{B(0, R)}), d_{\mathcal{W}_2} \right) \rightarrow \left(\mathcal{P}(\overline{B(0, R)}), \|\cdot\|_{H^{-2}(B(0, R))} \right)$$

is a continuous embedding.

Proof. Without loss of generality, let $R = 1/4$, and write $B := B(0, R)$. Fix $\mu, \nu \in \mathcal{P}(\overline{B(0, R)})$. Let $\pi \in \Pi(\mu, \nu)$ be any transport plan between μ and ν . We estimate $\|\mu - \nu\|_{H^{-2}(B(0, R))}$ as follows:

$$\begin{aligned} \|\mu - \nu\|_{H^{-2}}^2 &= \int_B |S(\mu - \nu)(\mathbf{x})|^2 d\mathbf{x} \\ &= \int_B |\Phi * (\mu - \nu)(\mathbf{x})|^2 d\mathbf{x} \\ &= \int_B \left(\iint_{B \times B} \Phi(\mathbf{x} - \mathbf{y}) - \Phi(\mathbf{x} - \mathbf{z}) d\pi(\mathbf{y}, \mathbf{z}) \right)^2 d\mathbf{x} \\ &\leq \iint_{B \times B} \int_B (\Phi(\mathbf{x} - \mathbf{y}) - \Phi(\mathbf{x} - \mathbf{z}))^2 d\mathbf{x} d\pi(\mathbf{y}, \mathbf{z}) \\ &\leq 2 \iint_{B \times B} \int_{A_{\mathbf{y}, \mathbf{z}}} (\Phi(\mathbf{a}) - \Phi(\mathbf{a} + \mathbf{y} - \mathbf{z}))^2 d\mathbf{a} d\pi(\mathbf{y}, \mathbf{z}), \end{aligned}$$

where $A_{\mathbf{y}, \mathbf{z}} := \{\mathbf{a} \in B(0, 2R) : |\mathbf{a}| \leq |\mathbf{a} + \mathbf{y} - \mathbf{z}|\}$. Note that $0 \leq \ln |\mathbf{a} + \mathbf{y} - \mathbf{z}| - \ln |\mathbf{a}| \leq \frac{|\mathbf{y} - \mathbf{z}|}{|\mathbf{a}|}$ and that

$$|\ln |\mathbf{a}| - \ln |\mathbf{a} + \mathbf{y} - \mathbf{z}|| \leq |\ln |\mathbf{a}||.$$

Now let $\delta = \frac{1}{2}|\mathbf{y} - \mathbf{z}|$. Continuing our estimates, we find

$$\begin{aligned}
\|\mu - \nu\|_{H^{-2}}^2 &\lesssim \iint_{B \times B} \int_{A_{\mathbf{y}, \mathbf{z}}} (\ln |\mathbf{a}| - \ln |\mathbf{a} + \mathbf{y} - \mathbf{z}|)^2 d\mathbf{a} d\pi(\mathbf{y}, \mathbf{z}) \\
&\lesssim \iint_{B \times B} \left(\int_{B(0, \delta)} \ln^2 |\mathbf{a}| d\mathbf{a} + \int_{A_{\mathbf{y}, \mathbf{z}} \setminus B(0, \delta)} \frac{|\mathbf{y} - \mathbf{z}|^2}{|\mathbf{a}|^2} d\mathbf{a} \right) d\pi(\mathbf{y}, \mathbf{z}) \\
&\lesssim \iint_{B \times B} |\mathbf{y} - \mathbf{z}|^2 \ln^2 |\mathbf{y} - \mathbf{z}| + |\mathbf{y} - \mathbf{z}|^2 \ln |\mathbf{y} - \mathbf{z}| d\pi(\mathbf{y}, \mathbf{z}) \\
&\lesssim \iint_{B \times B} |\mathbf{y} - \mathbf{z}| d\pi(\mathbf{y}, \mathbf{z}) \\
&\lesssim \left(\iint_{B \times B} |\mathbf{y} - \mathbf{z}|^2 d\pi(\mathbf{y}, \mathbf{z}) \right)^{1/2}.
\end{aligned}$$

Now let $\bar{\pi}$ be the transport plan that minimizes the quadratic Wasserstein distance, and since the above estimates hold for all transport plans we deduce

$$\|\mu - \nu\|_{H^{-2}}^2 \lesssim \left(\iint_{B \times B} |\mathbf{y} - \mathbf{z}|^2 d\bar{\pi}(\mathbf{y}, \mathbf{z}) \right)^{1/2} = d_{\mathcal{W}_2}(\mu, \nu).$$

□

Corollary 41. *Let $Y \subseteq \mathbb{R}^2$ be an open set that admits a Green's function. Then the identity map*

$$\text{id} : (\mathcal{P}(Y), d_{\mathcal{W}_2}) \rightarrow (\mathcal{P}(Y), \|\cdot\|_{H^{-2}})$$

is a continuous embedding.

Proof. Let $R > 0$ be such that $Y \subseteq B(0, R)$. By [29, Theorem 3.2.12] we have

$$G_Y(\mathbf{x}, \mathbf{y}) \leq G_{B(0, R)}(\mathbf{x}, \mathbf{y})$$

for all $\mathbf{x}, \mathbf{y} \in Y$, where these are the Green's functions on Y and $B(0, R)$ respectively. Hence,

$$\|\mu - \nu\|_{H^{-2}(Y)} \leq \|\mu - \nu\|_{H^{-2}(B(0, R))} \leq c \sqrt{d_{\mathcal{W}_2}(\mu, \nu)}$$

by using the previous theorem. □

The following general corollary follows by using a partition of unity to break up the support of a measure into single coordinate charts. In each chart we have the necessary metric bounds from the previous corollary. Finally, the result holds since M is compact and so only a finite number of charts are in play.

Corollary 42. *Let (M, g) be a compact Riemannian 2-manifold. Then the identity map*

$$\text{id} : (\mathcal{P}(M), d_{\mathcal{W}_2}) \rightarrow (\mathcal{P}(M), \|\cdot\|_{H^{-2}})$$

is a continuous embedding.

Following [5, 43], we extend the previous results to the case of signed measures by defining

$$\mathcal{W}_2(\mu, \nu) := d_{\mathcal{W}_2}(\mu^+ + \nu^-, \nu^+ + \mu^-).$$

Note that, unlike the usual Wasserstein distance, \mathcal{W}_2 is not a distance. Nevertheless, it provides an appropriate similarity measure for modeling the merging of cone singularities.

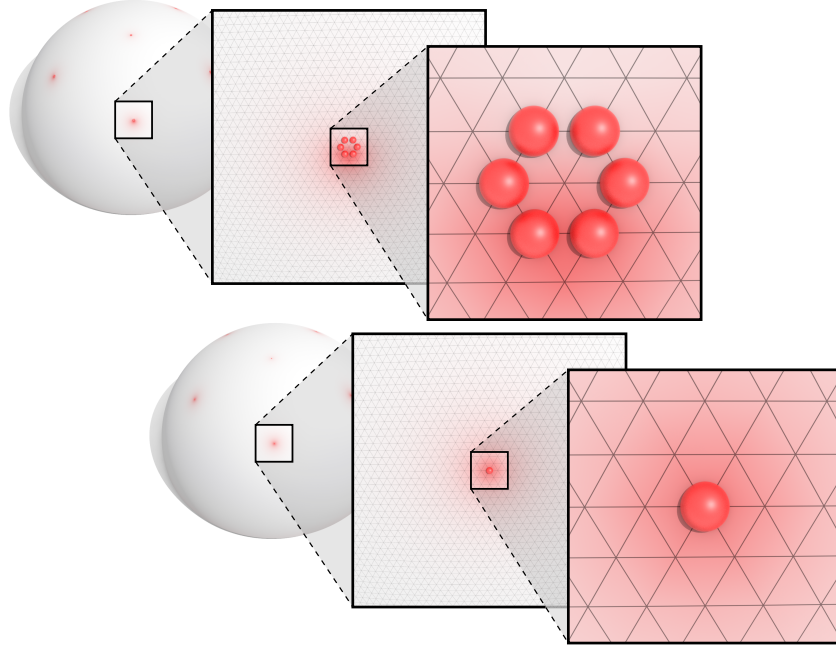


Figure 4.1: *Top:* for extremely fine meshes of very smooth surfaces (here, a hemisphere with 256k faces) our algorithm can produce cones arranged in tiny clusters rather than at isolated vertices. *Bottom:* A rigorous stability analysis shows that merging these cones to a nearby vertex (as shown here) cannot yield more than a miniscule change in area distortion, as indicated by red coloring. In practice, such merging is almost never needed.

The following result is a straight forward application of the Hahn-Banach theorem and the Riesz representation theorem for the dual space of continuous functions.

Lemma 43. *The vector space of linear combinations of Dirac deltas, $\text{Diracs}(M)$, supported at points in Y is dense in $\mathcal{M}(M)$ endowed the weak- * topology.*

Proof. By the Hahn-Banach theorem, to show that

$$\overline{\text{Diracs}(M)}^{\sigma(\mathcal{M}, \mathcal{C})} = \mathcal{M}(M)$$

it suffices to show that the only continuous (with respect to the weak- * topology) linear functional $L : \mathcal{M}(M) \rightarrow \mathbb{R}$ that satisfies $L(\mu) = 0$ for all $\mu \in \text{Diracs}(M)$ is given by $L \equiv 0$. By the Riesz representation theorem we know that every weak- * continuous linear functional L can be represented by a continuous function $\varphi \in \mathcal{C}(M)$ by

$$L(\mu) = \int_M \varphi d\mu.$$

Now consider such a linear functional annihilating $\text{Diracs}(M)$. Now for every $p \in M$ we have that

$$L(\delta_p) = \int_M \varphi d\delta_p = \varphi(p) = 0,$$

and so $\varphi \equiv 0$, which implies that $L \equiv 0$. Hence, $\text{Diracs}(M)$ is dense in $\mathcal{M}(M)$ with the weak- * topology. \square

We now can recover a well-known result with very little effort.

Corollary 44. *Let (M, g) be a compact 2-manifold. Then there exists a configuration of cone singularities such that the area distortion of the resulting conformal flattening is arbitrarily small.*

Proof. Fix $\varepsilon > 0$. Let $\Omega \in \mathcal{M}(M)$ denote the curvature 2-form. Note that $\mathcal{E}(u) = 0$ when $\Delta_g u = \Omega - \Omega$. Find $\delta > 0$ such that

$$|\mathcal{E}(u_\Omega) - \mathcal{E}_{L^2}(u_\mu)| < \varepsilon$$

for all $\mu \in \mathcal{M}(M)$ with $\mathcal{W}_2(\Omega, \mu) < \delta$. Here $\Delta u_\mu = \Omega - \mu$ and $\Delta u_\Omega = \Omega - \Omega = 0$, both with zero Dirichlet boundary conditions. We can find such a $\delta > 0$ using Theorem 40. By Lemma 43 there exists a sequence of finite linear combinations of delta measures $\{\mu_n\}_{n=1}^\infty$ such that $\mu_n \xrightarrow{*} \Omega$. By [43, Proposition 3.8] we deduce that

$$\mathcal{W}_2(\mu_n, \Omega) \rightarrow 0$$

since M is compact and metrizable. Take $N \in \mathbb{N}$ large enough such that $\mathcal{W}_2(\mu_N, \Omega) < \delta$, and write

$$\mu_N = \sum_{i=1}^m \alpha_i \delta_{p_i}, \quad \alpha_i \in \mathbb{R}, \quad p_i \in M.$$

It immediately follows that $\mathcal{E}(u_{\mu_N}) < \varepsilon$, as desired. \square

4.5 Moreau-Yosida Regularization

Numerically realizing optimization problems formulated over the non-reflexive space of measures may not be straightforward, and so we used the duality framework to obtain an equivalent optimization problem formulated over the usual function spaces. However, further numerical instabilities arise due to the extremely low regularity of the minimizers. In particular, if we try to directly use the optimality system to recover the minimizing measure we end up with serious numerical errors.

To derive a numerically stable and efficient algorithm we consider a sequence of regularized problems formulated over an intermediate Hilbert space. We begin by presenting this regularization in the context of the abstract framework, before specializing the results to our original cone parameterization problem. This regularized problem is introduced and analyzed in Clason & Kunisch [18]. Here, we present direct proofs based on elementary methods—such an exposition is clear with the abstract formulation of the problem.

4.5.1 Abstract Formulation

Let H be a Hilbert space such that $\mathcal{C} \hookrightarrow H$, where the embedding is dense. We identify $H \cong H^*$ using the Riesz representation theorem for Hilbert spaces. We write $(\mu, \nu)_H$ and $\|\mu\|_H$ to denote the inner product and norm in H . Assume that the inner product on H agrees with the duality pairing $\langle \cdot, \cdot \rangle_{\mathcal{M} \times \mathcal{C}}$, i.e.,

$$\langle \mu, \varphi \rangle_{\mathcal{M} \times \mathcal{C}} = (\mu, \varphi)_H, \quad \text{for all } \mu \in H, \quad \varphi \in \mathcal{C}.$$

Hence, we have the following *evolution triple*:

$$\mathcal{C} \hookrightarrow H \cong H^* \hookrightarrow \mathcal{M}.$$

Since \mathcal{C} is densely embedded in H we deduce that the adjoint map $H^* \hookrightarrow \mathcal{M}$ is injective. As is common in the study of evolutionary partial differential equations we will use the above evolution triple

(also known as rigged Hilbert space or as a Gelfand triple) to obtain better behaved optimization problems framed over the Hilbert space.

The regularized versions of (P_{abstract}) are parameterized by $\gamma > 0$:

$$\begin{array}{l} \text{minimize}_{\mu \in H, u \in U} \quad \mathcal{E}(u) + \mathcal{R}(\mu) + \frac{\gamma}{2} \|\mu\|_H^2 \\ \text{subject to} \quad s(u, \mu) = 0 \text{ in } X^*. \end{array} \quad (P_{\text{abstract}}^\gamma)$$

We also write (by a slight abuse of notation) $S : H \rightarrow U$ to denote the solution operator restricted to this Hilbert subspace.

To understand why the regularization formed over the Hilbert space is better behaved, note that for suitable choices of H and $s(\cdot, \cdot)$ the solution operator S will map H into a subspace (possibly a subset depending on the linearity of S) of much higher regularity than generic elements in $S(\mathcal{M})$. Our goal now is to discuss existence of solutions to $(P_{\text{abstract}}^\gamma)$ and show that solutions to this problem converge (in an appropriate topology) to minimizers of (P_{abstract}) .

Proposition 45. *The problem $(P_{\text{abstract}}^\gamma)$ admits a global solution $(\bar{\mu}_\gamma, S\bar{\mu}_\gamma) \in H \times U$.*

Proof. This is another direct application of the direct method. Since the cost functional contains the norm-squared, $\|\cdot\|_H^2$, it is clear that any minimizing sequence will be bounded in H . Hence, we can extract a weakly convergent subsequence in H (which also converges weakly- $*$ in \mathcal{M}). This limit is in fact a minimizer. \square

In a method that is completely analogous to the unregularized case we can obtain a necessary optimality system for $(P_{\text{abstract}}^\gamma)$. First we consider the reduced energy of $(P_{\text{abstract}}^\gamma)$

$$\begin{aligned} e_\gamma : H &\rightarrow \mathbb{R} \\ \mu &\mapsto e(\mu) + \frac{\gamma}{2} \|\mu\|_H^2 = \mathcal{E}(S\mu) + \frac{\gamma}{2} \|\mu\|_H^2. \end{aligned}$$

Remark. *For the optimization problem (P_{relaxed}) we will take $H = L^2(U_{\mathcal{R}})$. This choice is motivated by the relationship*

$$\|\mu\|_{\mathcal{M}(U_{\mathcal{R}})} = \|\mu\|_{L^1(U_{\mathcal{R}})}$$

for $\mu \in H$ (under an appropriate identification $H \subset \mathcal{M}$).

Proposition 46. *Let $\bar{\mu}_\gamma \in H$ be a minimizer of $(P_{\text{abstract}}^\gamma)$. Then $-\nabla e_\gamma(\bar{\mu}_\gamma) \in \partial \mathcal{R}(\bar{\mu}_\gamma)$.*

The rest of the necessary optimality system from (P_{abstract}) stays the same. Since the duality pairing in $\mathcal{M} \cong \mathcal{C}^*$ is compatible with the inner product in H we can easily compute the gradient of e_γ : for $\varphi \in \mathcal{C}$ we compute

$$\begin{aligned} de_\gamma(\mu)[\varphi] &= de(\mu)[\varphi] + \gamma(\mu, \varphi)_H \\ &= \langle \nabla e(\mu), \varphi \rangle_{\mathcal{M} \times \mathcal{C}} + \gamma(\mu, \varphi)_H \\ &= (\nabla e(\mu) + \gamma\mu, \varphi)_H. \end{aligned}$$

Therefore, $\nabla e_\gamma(\mu) = \nabla e(\mu) + \gamma\mu$.

Since we are dealing with a more regular optimization problem, we will be able to say more about the minimizers $\bar{\mu}_\gamma$. From Proposition 46 it immediately follows that

$$\bar{\mu}_\gamma = \arg \min_{\mu \in H} \left((\nabla e(\bar{\mu}_\gamma), \mu)_H + \frac{\gamma}{2} \|\mu\|_H^2 + \mathcal{R}(\mu) \right).$$

This expression justifies calling $(P_{\text{abstract}}^\gamma)$ a *Moreau-Yosida regularization*. Furthermore, even though this relationship includes $\bar{\mu}_\gamma$ on both sides of the equation it can be elegantly expressed using the *proximal map* of \mathcal{R} . Further discussion of the proximal map is postponed to Chapter 5.

4.5.2 Convergence of the Regularized Problems

Before we show that the regularized problems converge to the original problem as $\gamma \rightarrow 0^+$ we show that the *minimal cost* converges as $\gamma \rightarrow 0^+$. For notational convenience let's define

$$\begin{aligned} v_\gamma : H &\rightarrow \mathbb{R} \\ \mu &\mapsto v(\mu) + \frac{\gamma}{2} \|\mu\|_H^2 = e_\gamma(\mu) + \mathcal{R}(\mu) = \mathcal{E}(S\mu) + \mathcal{R}(\mu) + \frac{\gamma}{2} \|\mu\|_H^2. \end{aligned}$$

For $\gamma \geq 0$ define the optimal value function $o : [0, +\infty) \rightarrow \mathbb{R}$ by

$$o(\gamma) := v_\gamma(\bar{\mu}_\gamma),$$

where $\bar{\mu}_\gamma$ is a minimizer of $(P_{\text{abstract}}^\gamma)$. When $\gamma = 0$ we are referring to the quantities associated to the original unregularized problem. We can now obtain analogous results to [33, Theorem 4.1] through completely elementary methods regarding the concavity and differentiability of the optimal value functional:

Proposition 47. *The optimal value function $o : [0, +\infty) \rightarrow [v(\bar{\mu}), \infty)$ is concave and differentiable for (\mathcal{L}^1) -almost every $\gamma > 0$. Furthermore,*

$$o'(\gamma) = \frac{1}{2} \|\bar{\mu}_\gamma\|_H^2.$$

Proof. Fix $\gamma_0, \gamma_1 > 0$, and $\theta \in (0, 1)$. By minimality of $v_{\gamma_0}(\bar{\mu}_{\gamma_0})$ and $v_{\gamma_1}(\bar{\mu}_{\gamma_1})$ we have

$$\begin{aligned} \theta o(\gamma_0) + (1 - \theta) o(\gamma_1) &\leq \theta v_{\gamma_0}(\bar{\mu}_{\theta\gamma_0 + (1-\theta)\gamma_1}) + (1 - \theta) v_{\gamma_1}(\bar{\mu}_{\gamma_0}) \\ &= v(\bar{\mu}_{\theta\gamma_0 + (1-\theta)\gamma_1}) + \frac{\theta\gamma_0}{2} \|\bar{\mu}_{\theta\gamma_0 + (1-\theta)\gamma_1}\|_H^2 + \frac{(1-\theta)\gamma_1}{2} \|\bar{\mu}_{\theta\gamma_0 + (1-\theta)\gamma_1}\|_H^2 \\ &= v_{\theta\gamma_0 + (1-\theta)\gamma_1}(\bar{\mu}_{\theta\gamma_0 + (1-\theta)\gamma_1}) \\ &= o(\theta\gamma_0 + (1-\theta)\gamma_1). \end{aligned}$$

Thus, v is concave. Note that every concave function is (locally) Lipschitz, and so it is differentiable almost everywhere by Radamacher's theorem.

Now we compute the first derivative of o . For $\varepsilon > 0$ we see

$$\frac{o(\gamma_0 + \varepsilon) - o(\gamma_0)}{\varepsilon} \leq \frac{j_{\gamma_0 + \varepsilon}(\bar{\mu}_{\gamma_0}) - j_{\gamma_0}(\bar{\mu}_{\gamma_0})}{\varepsilon} = \frac{1}{2} \|\bar{\mu}_{\gamma_0}\|_H^2 \leq \frac{j_{\gamma_0}(\bar{\mu}_{\gamma_0}) - j_{\gamma_0 - \varepsilon}(\bar{\mu}_{\gamma_0})}{\varepsilon} \leq \frac{o(\gamma_0) - o(\gamma_0 - \varepsilon)}{\varepsilon}.$$

Taking $\varepsilon \rightarrow 0^+$ yields

$$\frac{d^+ o}{d\gamma}(\gamma_0) \leq \frac{1}{2} \|\bar{\mu}_{\gamma_0}\|_H^2 \leq \frac{d^- o}{d\gamma}(\gamma_0),$$

where $\frac{d^\pm o}{d\gamma}$ denote the directional derivatives of o with respect to γ . Thus, $o'(\gamma) = \frac{1}{2}\|\bar{\mu}_\gamma\|_H^2$. \square

To show convergence of the regularized functionals (the values of the minimizers, not necessarily the minimizers themselves) for $\gamma \rightarrow 0^+$ it suffices to show that o is continuous at zero. To deduce such convergence we only need the following additional assumption:

Assumption 6. For $\bar{\mu} \in \mathcal{M}$ there exists a sequence $\{\mu_n\}_{n=1}^\infty \subset H$ satisfying

$$\mu_n \xrightarrow{*} \bar{\mu} \text{ in } \mathcal{M} \quad \text{and} \quad \mathcal{R}(\mu_n) \rightarrow \mathcal{R}(\bar{\mu}) \quad \text{as } n \rightarrow +\infty.$$

Proposition 48. *The optimal value function o is continuous. In particular, $o(0) = \lim_{\gamma \rightarrow 0^+} o(\gamma)$.*

Proof. Fix any minimizer of the original problem $\bar{\mu} \in \mathcal{M}$ and consider a sequence $\{\mu_n\}_{n=1}^\infty \subset H$ such that $\mu_n \xrightarrow{*} \bar{\mu}$ in \mathcal{M} and $\mathcal{R}(\mu_n) \rightarrow \mathcal{R}(\bar{\mu})$. Then for all $\gamma > 0$ and $n \in \mathbb{N}$ we compute

$$\begin{aligned} o(0) &= v(\bar{\mu}) \leq v(\bar{\mu}_\gamma) \leq v(\bar{\mu}_\gamma) + \frac{\gamma}{2}\|\bar{\mu}_\gamma\|_H^2 \\ &= o(\gamma) \\ &\leq v(\mu_n) + \frac{\gamma}{2}\|\mu_n\|_H^2. \end{aligned} \tag{4.4}$$

By the convergence of $\mu_n \xrightarrow{*} \bar{\mu}$ and $\mathcal{R}(\mu_n) \rightarrow \mathcal{R}(\bar{\mu})$ we have

$$v(\mu_n) = e(\mu_n) + \mathcal{R}(\mu_n) \rightarrow e(\bar{\mu}) + \mathcal{R}(\bar{\mu}) = v(\bar{\mu})$$

as $n \rightarrow +\infty$. Now fix $\varepsilon > 0$ and take $N \gg 1$ large enough such that

$$v(\bar{\mu}) \leq v(\mu_N) \leq v(\bar{\mu}) + \frac{\varepsilon}{2}.$$

Using this in Eqn. 4.4 we obtain

$$o(0) \leq o(\gamma) \leq v(\mu_N) + \frac{\gamma}{2}\|\mu_N\|_H^2 \leq v(\bar{\mu}) + \frac{\varepsilon}{2} + \frac{\gamma}{2}\|\mu_N\|_H^2.$$

Since this inequality holds for all $\gamma > 0$, we can take $\gamma_\varepsilon \leq \varepsilon/\|\mu_N\|_H^2$ to obtain that

$$v(\bar{\mu}) \leq o(\gamma_\varepsilon) \leq v(\bar{\mu}) + \varepsilon.$$

Taking $\varepsilon \rightarrow 0^+$ shows the continuity of o at zero. \square

This continuity implies (up to a subsequence) the weak- $*$ convergence of the regularized minimizers $\bar{\mu}_\gamma$ to $\bar{\mu}$ as $\gamma \rightarrow 0^+$.

Theorem 49. *Consider a sequence $\gamma_n \rightarrow 0$ as $n \rightarrow +\infty$, and let $\bar{\mu}_{\gamma_n} \in H$ be the minimizers of $(P_{\text{abstract}}^{\gamma_n})$. There exists a minimizer $\bar{\mu} \in \mathcal{M}$ of (P_{abstract}) and a subsequence of γ_n (not relabeled) such that*

$$\bar{\mu}_{\gamma_n} \xrightarrow{*} \bar{\mu}$$

in \mathcal{M} as $n \rightarrow +\infty$. If the problem (P_{abstract}) has a unique solution $\bar{\mu} \in \mathcal{M}$ then $\bar{\mu}_\gamma \xrightarrow{} \bar{\mu}$ for any sequence $\gamma \rightarrow 0^+$.*

Proof. Note that the values of $\mathcal{R}(\bar{\mu}_{\gamma_n})$ are bounded by

$$\mathcal{R}(\bar{\mu}_{\gamma_n}) \leq v_{\gamma_n}(\bar{\mu}_{\gamma_n}) \leq \mathcal{E}(S\bar{\mu}) + \mathcal{R}(\bar{\mu})$$

due to the optimality of $v_{\gamma_n}(\bar{\mu}_{\gamma_n})$ for any fixed $\tilde{\mu} \in \mathcal{M}$. Since \mathcal{R} is coercive, $\{\mu_{\gamma_n}\}$ is bounded. Hence, we can extract a subsequence of γ_n (again, not relabeled) such that $\bar{\mu}_{\gamma_n} \xrightarrow{*} \hat{\mu}$ for some $\hat{\mu} \in \mathcal{M}$. Due to the weak- $*$ sequential lower semicontinuity of v it follows that

$$v(\hat{\mu}) \leq \liminf_{n \rightarrow +\infty} v_{\gamma_n}(\bar{\mu}_{\gamma_n}) = o(0) = v(\bar{\mu})$$

by Proposition 48. Hence, $\hat{\mu} \in \mathcal{M}$ is a minimizing solution of (P_{abstract}) .

Recall $\bar{\mu}_{\gamma_n} \xrightarrow{*} \bar{\mu}$ is equivalent to every subsequence having a further subsequence which converges to $\bar{\mu}$. If $\bar{\mu} \in \mathcal{M}$ is the unique minimizer then the convergence of the entire sequence $\gamma_n \rightarrow 0^+$ follows since we can apply the above result to an arbitrary subsequence $\{\bar{\mu}_{\gamma_{n_k}}\}_{k=1}^{\infty} \subseteq \{\bar{\mu}_{\gamma_n}\}_{n=1}^{\infty}$. Since this holds for every sequence $\gamma_n \rightarrow 0$ as $n \rightarrow +\infty$ we conclude that $\bar{\mu}_{\gamma} \xrightarrow{*} \bar{\mu}$ as $\gamma \rightarrow 0^+$. \square

4.5.3 Regularization for Optimal Cones

We finally return to our convex optimization problem (P_{relaxed}) . As noted before, we consider the Hilbert space $H = L^2(U_{\mathcal{R}}, dA)$. It obviously satisfies the assumption that $\mathcal{C}(U_{\mathcal{R}})$ is densely embedded in $L^2(U_{\mathcal{R}}, dA)$ and so the adjoint embedding $L^2(U_{\mathcal{R}}, dA) \hookrightarrow \mathcal{M}(U_{\mathcal{R}})$ is injective. Explicitly, the embedding is given by the map $\mu \in L^2(U_{\mathcal{R}}, dA) \mapsto \mu dA$.

Remark. Note that for the space $L^2(U_{\mathcal{R}}, dA)$ to be a meaningful (non-trivial) vector space we need to require that $U_{\mathcal{R}}$ is the closure of an open set. Furthermore, if we are to assume total boundary control then instead of considering $H = L^2(U_{\mathcal{R}}, dA)$ we consider the Hilbert space

$$H = L^2(U_{\mathcal{R}}, dA) \oplus L^2(\partial M, dl),$$

where dA and dl are the Riemannian volume forms on M and ∂M , respectively. It is necessary to make this distinction since the boundary ∂M is of measure zero with respect to the volume form on M , and so the L^2 -space would be completely agnostic to values on ∂M .

The regularized sequence of problems can be written as

$$\begin{array}{ll} \text{minimize} & \mathcal{E}(u) + \mathcal{R}(\mu) + \frac{\gamma}{2} \|\mu\|_{L^2(M)}^2 \\ \begin{array}{l} \mu \in \mathcal{M}(U_{\mathcal{R}}), \\ u \in W^{1,p}(M) \end{array} & \\ \text{subject to} & \Delta u = \Omega - \mu. \end{array} \quad (P_{\text{relaxed}}^{\gamma})$$

Recall that the area distortion energy

$$\mathcal{E}(u) := \int_{U_{\mathcal{R}}} |u|^2 w_{\mathcal{E}} dA$$

is convex. Since the solution operator for the singular Yamabe equation is affine linear we conclude that the reduced value functional $e : \mathcal{M} \rightarrow \mathbb{R}$ given by $e = \mathcal{E} \circ S_{\Omega}$ is also convex. Thus, we conclude that $(P_{\text{relaxed}}^{\gamma})$ admits a *unique* minimizer since the reduced value functional

$$v_{\gamma}(\cdot) = v(\cdot) + \frac{\gamma}{2} \|\cdot\|_H^2$$

is strongly convex.

Proposition 50. *There is a unique solution to $(P_{\text{relaxed}}^{\gamma})$.*

To ensure that we have convergence of the regularized problems to (P_{relaxed}) we need to verify Assumption 6 in this setting.

Proposition 51 (Assumption 6.). *Let $U_1 \subset M$ be an open subset and $U_2 \subset \partial M$ be an open subset in the induced topology on ∂M . Set $U := U_1 \cup U_2$. Let $\mathcal{R} : \mathcal{M}(U) \rightarrow \mathbb{R}$ be given by*

$$\mathcal{R}(\mu) = \int_U w(\mathbf{x}) d|\mu|(\mathbf{x})$$

for a positive weight function $w : U \rightarrow \mathbb{R}_{>0}$ that is bounded away from zero. Then for every $\mu \in \mathcal{M}(U)$ there exists a sequence $\{\mu_n\}_{n=1}^\infty \subset \mathcal{M}(U)$ such that $\mu_n \xrightarrow{} \mu$ in $\mathcal{M}(U)$ as $n \rightarrow +\infty$ and $\mathcal{R}(\mu_n) \rightarrow \mathcal{R}(\mu)$.*

Remark. *This result appears as Problem 24(B) in Brezis [12]. The proof is based on a clever application of the Hahn-Banach theorem. Alternatively, one can prove this directly using convolution of measures.*

In light of the above proposition and Theorem 49 we deduce

Theorem 52. *Let $\bar{\mu}_\gamma \in H$ be a sequence of L^2 -functions minimizing $(P_{\text{relaxed}}^\gamma)$ with regularization parameter γ . Similarly, let $\bar{\mu} \in \mathcal{M}(U_{\mathcal{R}})$ minimize (P_{relaxed}) . Then $\bar{\mu}_\gamma \xrightarrow{*} \bar{\mu}$ in $\mathcal{M}(U_{\mathcal{R}})$ as $\gamma \rightarrow 0^+$.*

Remark. *Note that by elliptic regularity of the Laplace-Beltrami operator since $\bar{\mu}_\gamma \in L^2(M, dA)$ we have $\bar{u}_\gamma \in H^2(M) \cap H_0^1(M)$ (or with other appropriate boundary conditions).*

We briefly recall the gradient of the reduced energy in the regularized case: if we let φ be a solution to the adjoint state equation $\Delta \varphi = w_\varepsilon u$, where $u = S\mu$, then $\nabla e_\gamma(\mu) = -\varphi + \gamma\mu$. We can summarize the entire first order optimality system as follows:

Proposition 53. *Let $\bar{\mu}_\gamma \in L^2(U_{\mathcal{R}})$ denote the minimizer of $(P_{\text{relaxed}}^\gamma)$. Let \bar{u}_γ denote the corresponding log-conformal scale factors. Then there exists an adjoint state $\bar{\varphi}_\gamma \in \mathcal{C}(U_{\mathcal{R}})$ such that the following regularized optimality system holds:*

$$\begin{cases} \Delta_g \bar{u}_\gamma = K - \bar{\mu}_\gamma, \\ \Delta_g \bar{\varphi}_\gamma = w_\varepsilon \bar{u}_\gamma, \\ \bar{\varphi}_\gamma - \gamma \bar{\mu}_\gamma \in \partial \mathcal{R}(\bar{\mu}_\gamma). \end{cases} \quad (\text{OS}_{\text{cones}}^\gamma)$$

Using the specific structure of our optimization problem we can explicitly characterize the error induced by the regularization. For simplicity, we consider $U_\varepsilon = U_{\mathcal{R}} = M$, $w_\varepsilon = 1$, and $w_{\mathcal{R}} = \lambda > 0$ —although these restrictions can be easily generalized without much work.

Proposition 54. *The regularization error can be written as*

$$o(\gamma) - o(0) = \frac{1}{2} \int_0^\gamma \|\bar{\mu}_\eta\|_H^2 d\eta$$

Proof. Since o is locally Lipschitz continuous it is locally absolutely continuous. Hence, by the fundamental theorem of calculus we see that

$$o(\gamma) - o(0) = \int_0^\gamma o'(\eta) d\eta.$$

By Proposition 47 we see that $o'(\eta) = \frac{1}{2} \|\bar{\mu}_\eta\|_H^2$ and so we obtain the desired result by substitution. \square

It seems likely that we can use Hölder bounds on the adjoint state to show that the above error estimate is on the order of γ^θ for some $0 < \theta < 1$. Unfortunately, I found a mistake in my original “proof” of this fact. A sketch of the argument is provided below:

The following result is a straightforward application of the variational inequality characterizing the minimizer and the definition of the proximal map—it will be shown explicitly in Chapter 5.

Proposition 55. *For $\gamma > 0$ we have that*

$$\bar{\mu}_\gamma = \frac{1}{\gamma} \max(0, \bar{\varphi}_\gamma - \lambda) + \frac{1}{\gamma} \min(0, \bar{\varphi}_\gamma + \lambda),$$

where $\bar{\varphi}_\gamma$ is the corresponding optimal adjoint state.

Proposition 56. *For $\gamma > 0$ we have that $o(\gamma) - o(0) < \gamma^\theta$ for some $\theta \in (0, 1)$.*

Proof. This is only a sketch, with only one crucially missing component!

Let $\gamma > 0$. By Hölder's inequality

$$\|\bar{\mu}_\gamma\|_{L^2(M)}^2 \leq \|\bar{\mu}_\gamma\|_{L^1(M)} \|\bar{\mu}_\gamma\|_{L^\infty(M)}.$$

Since $\|\bar{\mu}_\gamma\|_{L^1(M)} = \|\bar{\mu}_\gamma\|_{\mathcal{M}(M)}$ where we use the adjoint embedding $L^2(M) \cong L^2(M)^* \hookrightarrow \mathcal{M}(M)$ to identify $\bar{\mu}_\gamma$ with the appropriate measure, we see that

$$\|\bar{\mu}_\gamma\|_{L^1(M)} \leq o(\gamma) \leq \mathcal{E}(S_\Omega(0)),$$

since $\mathcal{R}(0) = 0$. So we see that

$$\|\bar{\mu}_\gamma\|_{L^2(M)}^2 \leq c \|\bar{\mu}_\gamma\|_{L^\infty(M)}$$

for some universal constant $c > 0$. Thus, using Proposition 54 we obtain the estimate

$$o(\gamma) - o(0) = \frac{1}{2} \int_0^\gamma \|\bar{\mu}_\eta\|_{L^2(M)}^2 d\eta \leq \frac{c}{2} \int_0^\gamma \|\bar{\mu}_\eta\|_{L^\infty(M)} d\eta.$$

So we see that we need an estimate of the form $\|\bar{\mu}_\eta\|_{L^\infty(M)} \leq c\eta^\beta$ for some $\beta > -1$. Using the elliptic regularity of the Laplace-Beltrami operator with measure-valued right hand sides we see that $\|\bar{\mu}_\eta\|_{W^{1,p}(M)} \leq c\|\bar{\mu}_\eta\|_{L^1(M)}$. Similarly, we conclude by using classical L^p -elliptic regularity and the Morrey embedding $\|\bar{\varphi}_\eta\|_{L^\infty} \leq \|\bar{\varphi}_\eta\|_{W^{3,p}(M)} \leq \|\bar{\mu}_\eta\|_{L^1(M)} \leq c$. So using Proposition 55 we see that

$$\|\bar{\mu}_\eta\|_{L^\infty(M)} \leq \frac{1}{\gamma} (\|\bar{\varphi}_\gamma\|_{L^\infty(M)} + \lambda).$$

Notice that this estimate is too weak to give us a useful bound since $\frac{1}{\eta}$ is not integrable near $\eta = 0$. *This is exactly the estimate that still needs work to give us an appropriate error estimate—it seems likely that one can use the fact that the Morrey embedding gives us an embedding into the Hölder spaces, and then use an interpolation inequality relating the Hölder norm weighted L^∞ and L^1 -norms.* \square

Chapter 5

Discretization

We will discuss the algorithmic realization of our optimal cones algorithm using the techniques presented in Chapter 4.

There are a number of ways to use the duality presented in Chapter 4 to obtain the optimal measure. The first idea is to directly solve the primal problem since it is a standard quadratic program, and then use the extremality conditions to recover the optimal measure. Numerical experiments (using MOSEK to solve the primal quadratic program) recover the optimal scale factors within some reasonable error, but we find in practice that recovering the optimal measure by computing $\Delta_g \bar{u} = \Omega - \bar{\mu}$ induces some serious numerical error which makes the recovery of the optimal cone angles and positions difficult. This numerical instability arises from the extremely low-regularity of the optimal measure, *i.e.*, the fact that the minimizing measures tend to represent linear combinations of Dirac deltas.

The key insight from the optimal control community is that it is necessary to intertwine the discretization with the solution of the optimization problem to produce numerically stable efficient algorithms. We use the techniques developed in [34, 64] to directly satisfy the first order optimality conditions. The main challenge with this approach is to reformulate the subgradient inclusion from the optimality system in a way that is suitable for numerical optimization.

5.1 Preliminaries

Our algorithm takes as input any manifold triangle mesh $\mathcal{X} = (V, E, F)$ with boundary $B \subset V$ (possibly empty), and produces as output a set of *cone vertices* $c_1, \dots, c_k \in V$ and associated *cone angles* ϕ_1, \dots, ϕ_k . This data can then be used to compute a parameterization using existing algorithms; for examples in this thesis we use *boundary first flattening (BFF)* [49] to compute the conformal parameterizations. The user must also provide a parameter $\lambda > 0$ which influences the number of cones (see Sec. 6.1.1 for further discussion).

Discrete Laplacian We discretize the Laplace-Beltrami operator Δ via the usual *cotan operator* [42], represented as a matrix $L \in \mathbb{R}^{|V| \times |V|}$. Let $\theta_i^{jk} \in \mathbb{R}$ denote the interior angle at vertex i of triangle ijk . Then L has nonzero entries

$$L_{ij} = \begin{cases} -\frac{1}{2} \sum_{ijk} \cot \theta_k^{ij}, & i \neq j, \\ -\sum_{il} L_{il}, & i = j, \end{cases}$$

where in the first case the sum is taken over triangles ijk containing edge ij , and in the second term the sum is taken over edges il containing vertex i . This matrix effectively encodes zero Neumann boundary conditions. We also build a diagonal (lumped) *mass matrix* $M \in \mathbb{R}^{|V| \times |V|}$ encoding the barycentric dual areas:

$$M_{ii} = \frac{1}{3} \sum_{ijk} A_{ijk}. \tag{5.1}$$

Here A_{ijk} is the area of triangle ijk , and we sum over triangles ijk containing vertex i .

To encode various boundary conditions, we partition the discrete Laplace-Beltrami matrix L into blocks corresponding to interior vertices and boundary vertices, denoted by I and B , respectively:

$$L = \begin{bmatrix} L_{II} & L_{IB} \\ L_{BI} & L_{BB} \end{bmatrix}$$

Note that $L_{IB} = L_{BI}^\top$ since L is symmetric. In general, the Poisson equation $\Delta u = \mu$ with Neumann boundary conditions $\frac{\partial u}{\partial n} = \nu$ on ∂M can be expressed as

$$\begin{bmatrix} L_{II} & L_{IB} \\ L_{BI} & L_{BB} \end{bmatrix} \begin{bmatrix} u_I \\ u_B \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \nu \end{bmatrix} = M\mu,$$

where ν_i and μ_i denotes the integrated value of ν and μ along the barycentric dual cell associated to the vertex $i \in V$. Similarly, if we want to consider the Dirichlet boundary conditions $u = b$ on ∂M then we have

$$\begin{bmatrix} L_{II} & L_{IB} \\ L_{BI} & L_{BB} \end{bmatrix} \begin{bmatrix} u_I \\ b \end{bmatrix} = M\mu,$$

which can be rewritten as solving for u_I :

$$L_{II}u_I = M_I\mu_I - L_{IB}b.$$

Discrete Curvature The integrated Gaussian curvature (*i.e.*, the curvature 2-form) associated with an interior vertex $i \in V$ can be discretized via the usual *angle defect*

$$\Omega_i := 2\pi - \sum_{ijk} \theta_i^{jk}. \quad (5.2)$$

These values are encoded in a column vector $\Omega \in \mathbb{R}^{|V|}$. This definition is motivated since Ω_i denotes the deviation of the angle sum around $i \in V$ from the angle sum around a planar triangulation (*i.e.*, $\Omega_i = 0$.) One should crucially note that Ω_i does not correspond to the *pointwise* Gaussian curvature, but rather the integrated Gaussian curvature over a region around the vertex. Similarly, we can define the discrete geodesic curvature 1-form along the boundary as

$$k_i := \pi - \sum_{ijk \in F} \theta_i^{jk}.$$

5.2 Semismooth Newton Method

We now introduce the necessary optimization background to reformulate the optimality system as a semismooth operator equation, and rely on the existing theory to reason about the convergence of the employed method [64].

The main idea behind *semismoothness* is that the subgradient is not the appropriate generalization of the differential to infinite dimensions, especially in the context of numerical optimization. We mainly follow [63, 64], but we do not consider multi-valued differentials (which is reminiscent of the subdifferential and more broadly of monotone operator theory) since it is both not useful to us from a practical perspective, and since it complicates the notation.

Definition 57. Let X and Y be Banach spaces, and $F : X \rightarrow Y$. Furthermore, let $DF : X \rightarrow \mathcal{L}(X, Y)$. The

map F is DF -semismooth at $x \in X$ if

$$\lim_{\|z\|_X \rightarrow 0^+} \frac{\|F(x+z) - F(x) - DF(x+z)[z]\|_Y}{\|z\|_X} = 0.$$

DF is the semismooth differential of F at x .

This definition varies from the usual definition of the Fréchet differential since the candidate differential DF can vary as $z \rightarrow 0$.

Remark. A careful look at the definition of semismoothness above shows that every function $F : X \rightarrow Y$ is semismooth differentiable with respect to some operator $DF : X \rightarrow \mathcal{L}(X, Y)$. That is for a given $\bar{x} \in X$ we can simply define $DF_{\bar{x}}$ by the relationship:

$$F(\bar{x} + z) - F(\bar{x}) - DF_{\bar{x}}(\bar{x} + z)[z] = 0$$

for all $z \in X$ (this corresponds to reaching \bar{x} in a single Newton iteration). However, for this notion to be useful algorithmically we must be able to compute DF a-priori, without knowledge of the point \bar{x} .

We now reproduce the proof from Hintermüller et al. [32, Theorem 1.1] on the convergence of Newton's method for solving the semismooth operator equation $F(x) = 0$. The proof is both short and instructive.

Theorem 58. Let $F : X \rightarrow Y$, $DF : X \rightarrow \mathcal{L}(X, Y)$, and $\bar{x} \in X$ satisfy

- $F(\bar{x}) = 0$,
- F is DF -semismooth at \bar{x} ,
- the semismooth differential DF is invertible in a neighborhood U of \bar{x} (in X) with a uniform bound

$$\|DF(x)^{-1}\|_{\mathcal{L}(X, Y)} \leq C, \quad \text{for all } x \in U.$$

Then there exists an open set V containing \bar{x} such that for $x_0 \in V$ the Newton iterations defined by

$$x_{n+1} = x_n - DF(x_n)^{-1}F(x_n)$$

converge to $\lim_{n \rightarrow +\infty} x_n = \bar{x}$. Furthermore, the convergence is superlinear in the sense that

$$\|x_{n+1} - x_n\|_X \leq \alpha_n \|x_n - \bar{x}\|_X$$

for a sequence $\{\alpha_n\}_{n=1}^{\infty} \subset [0, +\infty)$ satisfying $\alpha_n \rightarrow 0$ as $n \rightarrow +\infty$.

Proof. By using the definition of the Newton iterations we obtain

$$x_{n+1} - \bar{x} = -DF(x_n)^{-1} [F(x_n) - F(\bar{x}) - DF(x_n)[x_n - \bar{x}]].$$

Using the uniform bound on the operator norm of $DF(x_n)^{-1}$ we deduce

$$\|x_{n+1} - \bar{x}\|_X \leq C \|F(x_n) - F(\bar{x}) - DF(x_n)[x_n - \bar{x}]\|_Y$$

for all $x_n \in \bar{x} + U$. Since F is DF -semismooth at \bar{x} , we can take x_0 close enough to \bar{x} such that there exists a $0 \leq \eta < 1$ satisfying

$$\|F(\bar{x}) - F(\bar{x} + z) - DF(\bar{x} + z)[z]\|_Y \leq \frac{\eta}{M} \|z\|_X$$

for all $z \in X$ with $\|z\|_X \leq \|x_0 - \bar{x}\|_X$. It then follows that $\|x_1 - \bar{x}\|_X \leq \eta \|x_0 - \bar{x}\|_X$. By induction we obtain $\|x_n - \bar{x}\|_X \leq \eta^n \|x_0 - \bar{x}\|_X$. Therefore, we have $x_n \rightarrow \bar{x}$ as $n \rightarrow +\infty$. The superlinear convergence follows since we take η arbitrarily small by taking n larger and larger. \square

We recall some basic properties of the semismooth differential.

Proposition 59. *Let $F : X \rightarrow Y$ be \mathcal{C}^1 -Fréchet differentiable. Then F is dF -semismooth.*

Proposition 60 (Chain rule, [64, Proposition 3.8]). *Let $U \subset X$ and $V \subset Y$ be open. Let $G : U \rightarrow Y$ be Lipschitz continuous near $x \in U$ and DG -semismooth at x . Furthermore, let $F : V \rightarrow Z$ be DF -semismooth at $y = G(x)$ with DF being bounded near y . Let $G(U) \subset V$ and consider the operator $H = F \circ G : U \subset X \rightarrow Z$. Then H is DH -semismooth at x with*

$$DH(z) = DF(G(z)) \circ DG(z).$$

5.2.1 Semismooth Reformulation of the Regularized Optimality System

In light of the discussion on semismoothness in the previous section, an algorithm for computing the minimizers of the original and regularized problems arises if we show that the optimality systems from (OS_{cones}) and (OS_{cones}^Y) admit reformulations as semismooth operator equations. We will be able to show that the regularized problem admits such a semismooth reformulation since it is framed over the Hilbert space $H = L^2(M, dA)$. Similarly, we will be able to show that the finite-dimensional approximation of the original problem also admits a suitable formulation as a semismooth operator equation.

For notational simplicity we set $U_{\mathcal{R}} = U_{\mathcal{E}} = M$. The main difficulty in describing (OS_{cones}^Y) as an operator equation is in reformulating the relationship

$$-\nabla e_{\gamma}(\bar{\mu}) \in \partial \mathcal{R}(\bar{\mu}).$$

properly. We will show that this relationship admits a suitable reformulation using the *proximal map*.

Definition 61. Consider $f : H \rightarrow H$ be a convex function on the Hilbert space H . The *proximal map* of f , denoted $\text{Prox}_f : H \rightarrow H$, is given by

$$\text{Prox}_f(h) := \arg \min_{\tilde{h} \in H} \left(\frac{1}{2} \|\tilde{h} - h\|_H^2 + f(\tilde{h}) \right).$$

We will also write $\text{Prox}_{\gamma} := \text{Prox}_{f/\gamma}$ when the function f is understood from context.

We now recall the following basic result relating to the proximal mapping.

Proposition 62. *Let H be a Hilbert space, let $f : H \rightarrow H$ be a convex function, and let $\gamma > 0$. The map $\text{Prox}_{\gamma} : H \rightarrow H$ is a well-defined, bounded (nonlinear) operator. Furthermore,*

$$\mu = \text{Prox}_{\gamma}(\varphi) \iff \gamma(\varphi - \mu) \in \partial f(\mu).$$

Now we explicitly characterize this subdifferential relationship using the adjoint state. Let $\mu \in L^2(M, dA)$ and set $u = S_{\Omega}\mu$. The adjoint state $\varphi \in W_0^{1,p'}(M) \hookrightarrow \mathcal{C}(M)$ solves $\Delta \varphi = w_{\mathcal{E}}u$. Now using the expression for the reduced gradient we see that

$$-\nabla e_{\gamma}(\mu) \in \partial \mathcal{R}(\mu) \iff \gamma \left(\frac{1}{\gamma} \varphi - \mu \right) \in \partial \mathcal{R}(\mu) \iff \mu = \text{Prox}_{\gamma} \left(\frac{\varphi}{\gamma} \right).$$

This allows us to replace the subdifferential inclusion by the *equation* above.

For the regularized problem we consider the (semismooth) operator

$$F_\gamma : W_0^{1,p}(M) \times W_0^{1,p'}(M) \times L^2(M, dA) \rightarrow W_0^{1,p}(M)^* \times W_0^{1,p'}(M)^* \times L^2(M, dA)$$

$$F_\gamma(u, \varphi, \mu) := \begin{bmatrix} \Delta u - \Omega - \mu \\ \Delta \varphi - w_\mathcal{R} u \\ \mu - \text{Prox}_\gamma(\varphi/\gamma), \end{bmatrix}$$

where the proximal map is associated to the regularization \mathcal{R} . It is clear that regularized optimality system is satisfied if and only if $F_\gamma(\bar{u}, \bar{\varphi}, \bar{\mu}) = 0$. Note that the state and adjoint equations are only satisfied at the final iteration of the Newton method, and they are not satisfied during any other state of the system.

Proposition 63. Consider $\mathcal{R} : L^2(M, dA) \rightarrow \mathbb{R}$. Then

$$\text{Prox}_\gamma(\varphi) = \max\left(0, \varphi - \frac{1}{\gamma} w_\mathcal{R}\right) + \min\left(0, \varphi + \frac{1}{\gamma} w_\mathcal{R}\right),$$

where the max and min operators are taken pointwise almost-everywhere.

Proof. Recall that under the identification $L^2(M, dA) \hookrightarrow \mathcal{M}(M)$.

$$\mathcal{R}(\mu) = \int_M w_\mathcal{R}(\mathbf{x}) |\mu(\mathbf{x})| dA(\mathbf{x})$$

for $\mu \in L^2(M, dA)$. To find $\mu = \text{Prox}_\gamma(\varphi)$ we need to minimize the functional

$$\mu \mapsto \int_M \frac{\gamma}{2} |\mu(\mathbf{x}) - \varphi(\mathbf{x})|^2 + w_\mathcal{R}(\mathbf{x}) |\mu(\mathbf{x})| dA(\mathbf{x}).$$

Clearly, it suffices to minimize the integrand pointwise (dA -almost everywhere). Fix $\mathbf{x} \in M$. Minimizing the one-dimensional function

$$m \mapsto \frac{\gamma}{2} (m - \varphi(\mathbf{x}))^2 + w_\mathcal{R}(\mathbf{x}) |m|$$

gives rise to

$$m = \begin{cases} 0 & \text{if } |\varphi(\mathbf{x})| \leq \frac{1}{\gamma} w_\mathcal{R}(\mathbf{x}), \\ \varphi(\mathbf{x}) - \frac{1}{\gamma} \text{sgn}(\varphi(\mathbf{x})) w_\mathcal{R}(\mathbf{x}) & \text{else.} \end{cases}$$

So we conclude that

$$\text{Prox}_\gamma(\varphi) = \max\left(0, \varphi - \frac{1}{\gamma} w_\mathcal{R}\right) + \min\left(0, \varphi + \frac{1}{\gamma} w_\mathcal{R}\right),$$

where max and min operators are understood in a pointwise (dA -almost-everywhere) sense. \square

Since we know that $\mu = \text{Prox}_\gamma(\varphi/\gamma)$ we can in fact reduce the optimality system for the regularized problem even further to consider $F_\gamma : W_0^{1,p}(M) \times W_0^{1,p'}(M) \rightarrow W_0^{1,p}(M)^* \times W_0^{1,p'}(M)^*$ (with a slight overloading of notation)

$$F_\gamma(u, \varphi) = \begin{bmatrix} \Delta u - \Omega - \text{Prox}_\gamma(\varphi/\gamma) \\ \Delta \varphi - w_\mathcal{R} u \end{bmatrix}$$

It is important to note that Prox_γ is semismooth differentiable as a map from $L^r(M) \rightarrow L^2(M)$ for any $r > 2$. Although this may initially seem problematic since the control $\mu \in L^2(M)$, it is not since by the elliptic regularity of the Laplacian and the Sobolev embedding theorems the gradient is a map

$\nabla e : L^2(M) \rightarrow L^1(M)$. Thus, the proximal map is semismooth differentiable (with respect to $D \text{Prox}_\gamma$ described below) in the image of ∇e . A direct computation shows that

$$D \text{Prox}_\gamma(\varphi/\gamma)[\psi] = \begin{cases} \psi(\mathbf{x})/\gamma & \text{if } |\varphi(\mathbf{x})| > w_{\mathcal{D}}(\mathbf{x}), \\ 0 & \text{else.} \end{cases}$$

By a slight abuse of notation we write $D \text{Prox}_\gamma(\varphi/\gamma) = \frac{1}{\gamma} \mathbb{1}_{|\varphi| > w_{\mathcal{D}}}$ since the application is understood as pointwise multiplication of this characteristic function.

Using the chain rule we deduce that F_γ is semismooth differentiable with

$$DF_\gamma(u, \varphi) = \begin{bmatrix} \Delta & -D \text{Prox}_\gamma(\varphi/\gamma) \\ -w_{\mathcal{D}} & \Delta \end{bmatrix} = \begin{bmatrix} \Delta & -\frac{1}{\gamma} \mathbb{1}_{|\varphi| > w_{\mathcal{D}}} \\ -w_{\mathcal{D}} & \Delta \end{bmatrix},$$

where the application of $DF_\gamma(u, \varphi)$ on tangent vectors (v, ψ) is given through the notation of matrix multiplication.

5.3 Algorithm

To numerically approximate solutions to the original and relaxed minimization problems we use the standard approach of finite elements. We only briefly sketch the main ideas here since we exactly follow the approach of Casas et al. [13]. We consider piecewise linear finite elements on the triangulated surface $\mathcal{X} = (V, E, F)$ – the basis elements are the Whitney hat functions:

$$\{\varphi_i \in \text{PL}(\mathcal{X}; \mathbb{R}) : i \in V, \varphi_i = 1, \varphi_j = 0 \text{ for all } j \neq i\}.$$

Now we have the finite dimensional approximation of the space $\mathcal{C}(M)$ given by

$$C_{\mathcal{X}} := \left\{ \sum_{i \in V} c_i \varphi_i : c_i \in \mathbb{R} \right\}.$$

Now we can define a finite dimensional space of measures as the dual of $X_{\mathcal{X}} := C_{\mathcal{X}}^*$ to obtain

$$X_{\mathcal{X}} := \left\{ \mu = \sum_{i \in V} \mu_i \delta_{v_i} : \mu_i \in \mathbb{R} \right\}.$$

Note that the duality pairing between $X_{\mathcal{X}}$ and $C_{\mathcal{X}}$ coincides with the natural duality pairing given by the inner product on $\mathbb{R}^{|V|}$. When we restrict the Laplace-Beltrami operator to this finite dimensional subspace we obtain that the Laplace-Beltrami operator is replaced by the cotan-Laplacian L mentioned in Sec. 5.1. Our optimization problems are now discretized in the straightforward way using these finite elements.

Note that we only need to use this (slightly non-standard) conforming discretization of the space of measures $\mathcal{M}(M)$ in the final stage to actually recover the original cone singularities. By first solving the sequence of regularized problems, which are framed over the Hilbert space $L^2(M)$, we are utilizing the much more well-established finite-element method for solving partial differential equations. Furthermore, attempts to directly solve the discretized optimization problem formulated over the space of measures results in many of the same issues that arise in the L^1 -regularized optimization. Understanding exactly the root of the differences between the L^1 -regularization and the measure-norm regularization in the discrete setting is a most natural and important next step.

Notation. We will use sans-serif fonts $u, p, \mu \in \mathbb{R}^{|V|}$ to denote the discrete values corresponding to the

smooth optimization variables u, φ, μ .

5.3.1 Semismooth Reformulation of the Discrete Optimality System

Since it will be important for implementation (and since the main ideas have been presented succinctly), we reintroduce the regions U_ε and $U_\mathcal{R}$ in this section. We let R and E be the diagonal matrices that indicate the domains $U_\mathcal{R}$ and U_ε respectively. Similarly, $W_\varepsilon := \text{diag}(w_\varepsilon)$. Note that one can consider $\mu \in \mathbb{R}^{|U_\mathcal{R}|}$, in which case we use a matrix $R \in \mathbb{R}^{|V| \times |U_\mathcal{R}|}$ to extend a vector on $U_\mathcal{R}$ to a vector on all of V by simply placing zeros at vertices in $V \setminus U_\mathcal{R}$.

Since we already worked out the semismooth reformulation for the regularized optimization problem in the continuous setting, we begin by considering the optimality system of the discrete *regularized* problems. In a manner that is completely parallel to the continuous derivation we deduce that for $\gamma > 0$ the discrete optimality system reads

$$\begin{cases} Lu = \Omega - R\mu, \\ L^\top p = EMW_\varepsilon u, \\ \mu = \frac{1}{\gamma} \text{Prox}(R^\top p). \end{cases} \quad (5.3)$$

Note that we can substitute the expression for μ into the first of these equations to obtain a system involving only the variables $u, p \in \mathbb{R}^{|V|}$. As in the unregularized case, we can describe this system as an equation $F_\gamma(u, p) = 0$ where $F_\gamma : \mathbb{R}^{2|V|} \rightarrow \mathbb{R}^{2|V|}$ is given by

$$F_\gamma(u, p) := \begin{bmatrix} Lu - \Omega + \frac{1}{\gamma} R \cdot \text{Prox}(R^\top p) \\ L^\top p - EMW_\varepsilon u \end{bmatrix}. \quad (5.4)$$

Note that this map is both nonlinear and nonsmooth due only to the proximal map. Compare this with the semismooth operator equation formulated for the regularized optimization problem framed over the original manifold, $(OS_{\text{cones}}^\gamma)$.

Now we turn to the original relaxed optimization problem. By applying either Fenchel-Rockafellar duality (or direct convex duality) to the discretized optimization we deduce that for the original problem the optimality system reads

$$\begin{cases} Lu = \Omega - R\mu, \\ L^\top p = EMW_\varepsilon u, \\ \mu \in \partial i_{\lambda w_\mathcal{R}}(R^\top p), \end{cases} \quad (5.5)$$

Since $\mathbb{R}^{|V|}$ is a Hilbert space, we will again be able to use the proximal map to characterize the above subdifferential inclusion. Note that such a reformulation was not possible over the non-reflexive Banach space $\mathcal{M}(M)$. Bauschke & Combettes [7] shows the following classical result regarding the relationship between the proximal map, the resolvent, and the subdifferential.

Definition 64. Let H be a Hilbert space, let $A : H \rightarrow 2^H$ be a multi-valued operator, and let $\gamma > 0$. The *resolvent* of A is defined to be

$$J_A = (\mathbb{1} + A)^{-1}.$$

The *Yosida approximation* of A of index γ is

$$\gamma A = \frac{1}{\gamma} (\mathbb{1} - J_{\gamma A}).$$

Proposition 65 (Bauschke & Combettes [7, Proposition 23.3]). *Let $A : H \rightarrow 2^H$, $\gamma > 0$, $x \in H$, and $p \in H$. Then the following hold:*

- (i) $\text{dom} J_{\gamma A} = \text{dom}^\gamma A = \text{range}(\mathbb{1} + \gamma A)$ and $\text{range} J_{\gamma A} = \text{dom} A$.
- (ii) $p \in J_{\gamma A}(x) \iff x \in p + \gamma A(p) \iff x - p \in \gamma A(p)$.
- (iii) $p \in {}^\gamma A(x) \iff p \in A(x - \gamma p)$.

We will use this result where $A = \partial i_{\lambda w_{\mathcal{R}}} : L^2(M) \rightarrow L^2(M)$ is the subdifferential of the indicator, which is a maximally monotone operator. It is not hard to see that

$$J_{\gamma} \partial f = \text{Prox}_{\gamma f}$$

for any convex $f : H \rightarrow H$ and $\gamma > 0$. Now if we take $f = i_{\lambda w_{\mathcal{R}}}$ then

$$J_{\gamma} \partial i_{\lambda w_{\mathcal{R}}} = \text{Prox}_{\gamma i_{\lambda w_{\mathcal{R}}}} = \text{Prox}_{i_{\lambda w_{\mathcal{R}}}} = \text{Proj}_{\lambda w_{\mathcal{R}}},$$

where $\text{Proj}_{\lambda w_{\mathcal{R}}}$ is the projection onto the box constraints $\lambda w_{\mathcal{R}}$. We can easily express this projection as:

$$\text{Proj}_{\lambda w_{\mathcal{R}}}(v) := v - \max(0, v - \lambda w_{\mathcal{R}}) - \min(0, v + \lambda w_{\mathcal{R}}).$$

Proposition 66. *The relationship $\mu \in \partial i_{\lambda w_{\mathcal{R}}}(\mathbf{q})$ is equivalent to, for all $\gamma > 0$,*

$$\mathbf{q} = \text{Proj}_{\lambda w_{\mathcal{R}}}\left(\mathbf{q} + \frac{\mu}{\gamma}\right).$$

Proof. By Proposition 65 (ii) we see that $\mu \in \partial i_{\lambda w_{\mathcal{R}}}(\mathbf{q})$ is equivalent to, for all $\gamma > 0$,

$$\mathbf{q} = J_{\gamma} \partial i_{\lambda w_{\mathcal{R}}}\left(\mathbf{q} + \frac{\mu}{\gamma}\right) = \text{Proj}_{\lambda w_{\mathcal{R}}}\left(\mathbf{q} + \frac{\mu}{\gamma}\right).$$

□

Utilizing this and the subdifferential relationship from the discrete optimality system we see that for all $\gamma > 0$

$$\mu \in \partial i_{\lambda w_{\mathcal{R}}}(\mathbf{R}^\top \mathbf{p}) \iff \mathbf{R}^\top \mathbf{p} = \text{Proj}_{i_{\lambda w_{\mathcal{R}}}}\left(\gamma \left(\frac{1}{\gamma} \mathbf{R}^\top \mathbf{p} + \mu\right)\right).$$

Expanding this out, we conclude that $\mu \in \partial i_{\lambda w_{\mathcal{R}}}(\mathbf{E}^\top \mathbf{p})$ is equivalent to

$$\mu = \gamma \max\left(0, \mathbf{E}^\top \mathbf{p} + \frac{\mu}{\gamma} - \lambda w_{\mathcal{R}}\right) + \gamma \min\left(0, \mathbf{E}^\top \mathbf{p} + \frac{\mu}{\gamma} + \lambda w_{\mathcal{R}}\right)$$

for all $\gamma > 0$. In particular, this holds for $\gamma = 1$ and so we use this relationship

$$\mu = \max\left(0, \mathbf{E}^\top \mathbf{p} + \mu - \lambda w_{\mathcal{R}}\right) + \min\left(0, \mathbf{E}^\top \mathbf{p} + \mu + \lambda w_{\mathcal{R}}\right)$$

to encode the discrete optimality system in the operator $F : \mathbb{R}^{3|V|} \rightarrow \mathbb{R}^{3|V|}$

$$F(\mathbf{u}, \mathbf{p}, \mu) := \begin{bmatrix} \mathbf{L}\mathbf{u} - \Omega - \mathbf{E}\mu \\ \mathbf{L}^\top \mathbf{p} - \mathbf{R}\mathbf{M}w_{\mathcal{R}}\mathbf{u} \\ \mu - \text{Proj}\left(\mathbf{E}^\top \mathbf{p} + \mu\right) \end{bmatrix}. \quad (5.6)$$

This is a valid reformulation since the optimality system holds if $F(\mathbf{u}, \mathbf{p}, \mu) = 0$.

Now we compute the semismooth differentials. By considering the subgradient it is straightforward to see that Proj is semismooth differentiable, with a semismooth differential given by the diagonal

matrix

$$D \text{Proj}(\mathbf{x})_{ii} := |\mathbf{x}_i| > \lambda w_{\mathcal{X}_i}.$$

The analogou computations were provided for the smooth regularized problem. We will write $D \text{Proj}(\mathbf{x}) = \mathbb{1}_{|\mathbf{x}| > \lambda w_{\mathcal{X}}}$. Using this, and the fact that the rest of the operators in F and F_γ are linear, we obtain that

$$DF_\gamma(\mathbf{u}, \mathbf{p}) := \begin{bmatrix} \mathbf{L} & \frac{1}{\gamma} \mathbf{R} \mathbb{1}_{|\mathbf{R}^\top \mathbf{p}| > \lambda w_{\mathcal{X}}} \\ -\mathbf{E} \mathbf{M} \mathbf{W}_{\mathcal{E}} & \mathbf{L}^\top \end{bmatrix} \quad (5.7)$$

and

$$DF(\mathbf{u}, \mathbf{p}, \mu) := \begin{bmatrix} \mathbf{L} & \mathbf{0} & \mathbf{R} \\ -\mathbf{E} \mathbf{M} \mathbf{W}_{\mathcal{E}} & \mathbf{L}^\top & \mathbf{0} \\ \mathbf{0} & \mathbb{1}_{|\mathbf{R}^\top \mathbf{p} + \mu| > \lambda w_{\mathcal{X}}} & \mathbb{1} - \mathbb{1}_{|\mathbf{R}^\top \mathbf{p}| > \lambda w_{\mathcal{X}}} \end{bmatrix} \quad (5.8)$$

In the above $\mathbb{1} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$. We again emphasize that these descriptions of the optimality system allow us to easily apply a semismooth Newton method to solve the equation $F(\mathbf{u}, \mathbf{p}, \mu) = 0$ and $F_\gamma(\mathbf{u}, \mathbf{p}) = 0$, as described in Sec. 5.2.

5.3.2 Implementation

With everything in place we can describe the final algorithm which we implement. The basic idea of the algorithm is to solve the first-order optimality conditions described in (OS_{cones}) . A standard practice for our particular type of problem is to solve it in two stages:

- **STAGE I** — Solve a sequence of regularized problems that approach the exact problem.
- **STAGE II** — Use the solution to the final regularized problem to initialize a solve for the exact solution.

We reiterate, the reason for solving a sequence of regularized problems first rather than directly solving the exact problem is that the original problem has extremely low regularity solutions and so without a good initialization the Newton method described can be unnecessarily slow and may not even converge; indeed, in practice we find that solving the regularized problem first speeds up computation by about an order of magnitude. In terms of our original geometric problem, **STAGE I** produces a sequence of sharper and sharper approximations of the optimal cone distribution (starting with a highly “smoothed out” version); the second stage is then used to recover the exact cones.

5.3.3 Active Sets

Note that the semismooth differential depends on the subset of M where either the max and min operators in Proj are non-zero.

These regions are called the *active sets*. For the regularized problem determined by F_γ these read

$$\begin{aligned} A^+ &:= \{\mathbf{x} \in M : p(\mathbf{x}) > +\lambda\}, \\ A^- &:= \{\mathbf{x} \in M : p(\mathbf{x}) < -\lambda\}. \end{aligned}$$

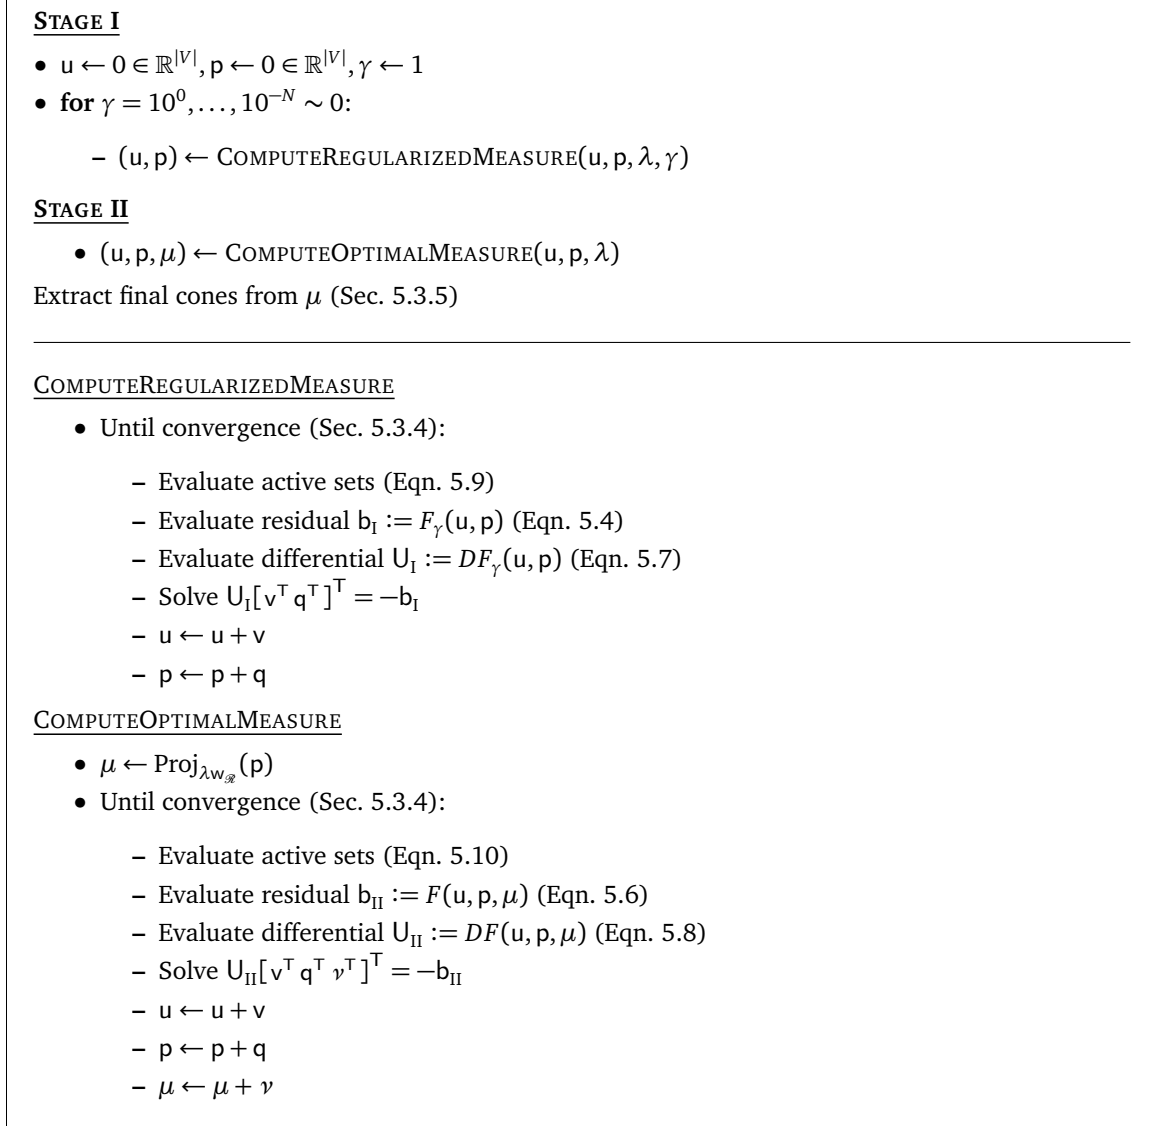


Figure 5.1: Our algorithm boils down to solving a sequence of sparse linear equations, together with some simple closed-form evaluations. In practice we use $N = 10$ for the largest regularization parameter.

Discretely, we consider the active sets $A^+, A^- : V \rightarrow \{0, 1\}$. We will represent these sets as $|V| \times |V|$ diagonal matrices; in **STAGE I**, these matrices have nonzero entries

$$\begin{aligned} (A_I^+)_{ii} &:= p_i > +\lambda, \\ (A_I^-)_{ii} &:= p_i < -\lambda \end{aligned} \tag{5.9}$$

for each vertex $i \in V$. Similarly, in **STAGE II**, where we also have variables $\mu \in \mathbb{R}^{|V|}$ the active sets are given by

$$\begin{aligned} (A_{II}^+)_{ii} &:= p_i + \mu_i > +\lambda, \\ (A_{II}^-)_{ii} &:= p_i + \mu_i < -\lambda, \end{aligned} \tag{5.10}$$

i.e., the vertices i where $|\rho_i + \mu_i| > \lambda$. These sets are needed to check for convergence and to build our update matrices, as described in Sections 5.3.1 and 5.3.4. In both stages, we will also define the matrix

$$A := A^+ + A^-. \quad (5.11)$$

5.3.4 Convergence Criteria

Let $(A^+)_k$ and $(A^-)_k$ denote the two active sets at the k th iteration of Newton's method (whether in **STAGE I** or **STAGE II**). To check convergence, we simply check if either

$$\begin{aligned} (A^+)_k &= (A^+)_{k-1}, \quad \text{and} \\ (A^-)_k &= (A^-)_{k-1}, \end{aligned}$$

indicating that the active sets are no longer changing, or if the norm of the current residual b is below a small tolerance $\varepsilon > 0$. An analysis of this stopping criteria is found in [32].

5.3.5 Extracting Cones

At the end of **STAGE II** we have a value μ at each vertex. To extract the final cones, we simply identify the vertices $c_i \in V$ where $\mu_{c_i} \neq 0$. Numerically, this is very easy to do since the values are extremely stratified, *i.e.*, they are either numerically zero (we use a tolerance of 10^{-12}), or they are equal to the cone angle ϕ_{c_i} . Very rarely (*e.g.*, only on extremely fine meshes) cones may appear in tiny clusters, reflecting the fact that in the smooth setting one can always obtain slightly smaller area distortion by replacing a Dirac measure at a point p with a measure supported on a tiny ring or region around p (see Sec. 4.4). In practice we simply replace each edge-connected set of cones c_{i_1}, \dots, c_{i_m} with a single cone of same total magnitude $\phi_{i_1} + \dots + \phi_{i_m}$ at the location of the Fréchet mean of these points, c_{mean} . Our Wasserstein stability result shows that this rounding procedure cannot change the resulting area distortion by more than a tiny amount related to the distance to c_{mean} (Sec. 4.4). See Fig. 4.1 for one example.

5.3.6 Multiresolution

Since the initial phase of **STAGE I** involves problems that are highly spatially regularized (*i.e.*, when γ is small), it makes little sense to solve these problems on a fine triangulation, where the scale of features in the optimal solution will be much larger than the typical edge length. Moreover, since the solutions to these problems are used only to initialize the next problem in the sequence, they do not need to be solved with high spatial accuracy. In practice we therefore adopt a simple multiresolution strategy: given our input mesh \mathcal{K} we first construct a sequence of progressively finer meshes $\mathcal{K}_1, \dots, \mathcal{K}_M = \mathcal{K}$, where M is no greater than the number N of outer iterations used in **STAGE I**, and the number of triangles in consecutive meshes is related by roughly the same constant factor s . The solution is then transferred from coarse to fine as the value of γ increases (in practice, we use each mesh roughly the same number of times). In particular, for each vertex i on a mesh \mathcal{K}_l , we identify the set of vertices on the next finest mesh \mathcal{K}_{l+1} that are closer to i than any other vertex in \mathcal{K}_l . The values of u_i and ρ_i are then equally distributed over these vertices. Note that one does not need to be particularly careful about the number of meshes used, nor the method of coarsening. Since our problem is convex, we will always find the same solution: the multiresolution strategy affects only the computational cost. In practice we use the *Reduce* functionality in *MeshMixer* [50], and use a constant $s = 2$. Fig. 5.2 shows one

example, where we obtain a speedup of roughly 30x. Note however that the method is still quite efficient even without this multiresolution strategy; most examples in this thesis were computed directly on the fine input mesh.




Figure 5.2: Mesh hierarchy built on the *Old Man Multires* model; here we use only three levels. Compared to solving directly on the full-resolution 30k triangle mesh, we obtain a speedup factor of about 23x (from 14.77s to about 0.62s). Colors indicate the optimal regularized measure; in the final (fine) mesh the measure is exactly concentrated onto isolated vertices. *Far right:* The resulting parameterization has only 32 cones, with a total angle of 18.86π .

Chapter 6

Results

6.1 Validation and Comparisons

We implemented our algorithm in C++ using double precision for all calculations and the sparse QR solver in *SuiteSparse* to solve linear systems. Timings were measured on a 2.6GHz Intel Core i5 laptop with 8GB of RAM. In practice we need to solve about 50 to 100 linear systems, independent of the type of geometry or the resolution of the model. The multiresolution strategy outlined in Sec. 5.3.6 reduces the size of these systems substantially, though we did not find it essential for most of the examples in this thesis: for models of about 100–150k triangles the algorithm takes at most about 20–25 seconds. We did very little optimization of our code, and there are plenty of opportunities for acceleration both at the level of linear algebra, and in terms of algorithms (as mentioned at the beginning of Sec. 5.3).

6.1.1 Tuning Parameter

The parameter $\lambda \geq 0$ influences the number of cone singularities, or more precisely, the maximum allowable total cone angle $\Phi = \sum_i |\phi_i|$. In particular, increasing λ decreases the total, and decreasing λ increases it. A greater total cone angle will result in lower area distortion—however, for values of λ that are too close to zero, cones will be placed everywhere. From basic scaling considerations and practical experience, we find that a good range of values across a wide variety of examples at different resolutions is $\frac{1}{10} \text{Area}(\mathcal{K}) < \lambda < \text{Area}(\mathcal{K})$ where $\text{Area}(\mathcal{K})$ is the total surface area (see Figures 6.2 and 6.1). For values below this range one typically starts to see cones densely distributed in regions rather than at isolated points; above this range one tends to get no cones at all. In the region $0 < \lambda < \frac{1}{10} \text{Area}(\mathcal{K})$ there are some choices that lead to configurations of cone singularities supported on curves. (See Fig. 6.1.)

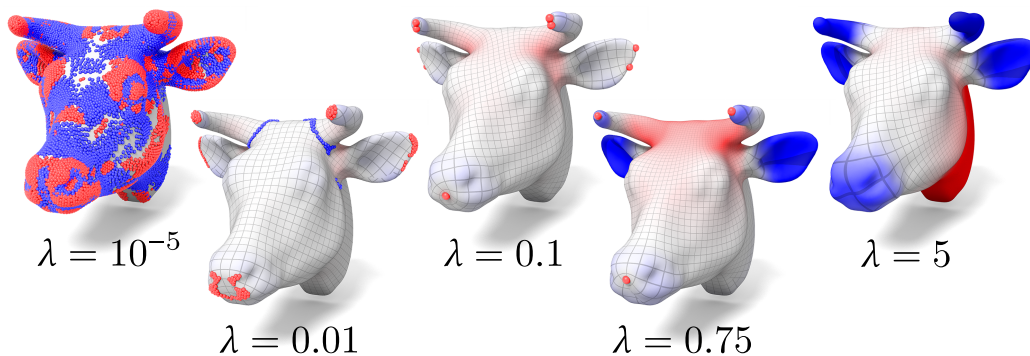


Figure 6.1: The parameter λ controls the strength of the penalty on the total cone angle, and in turn, influences the number of cones. For values above 1 (strong penalty) we tend to see that no cones are placed. For values very close to zero (weak penalty) the curvature measure stops being sparse, and we get cones with small angles placed densely across regions or along curves rather than at isolated points.

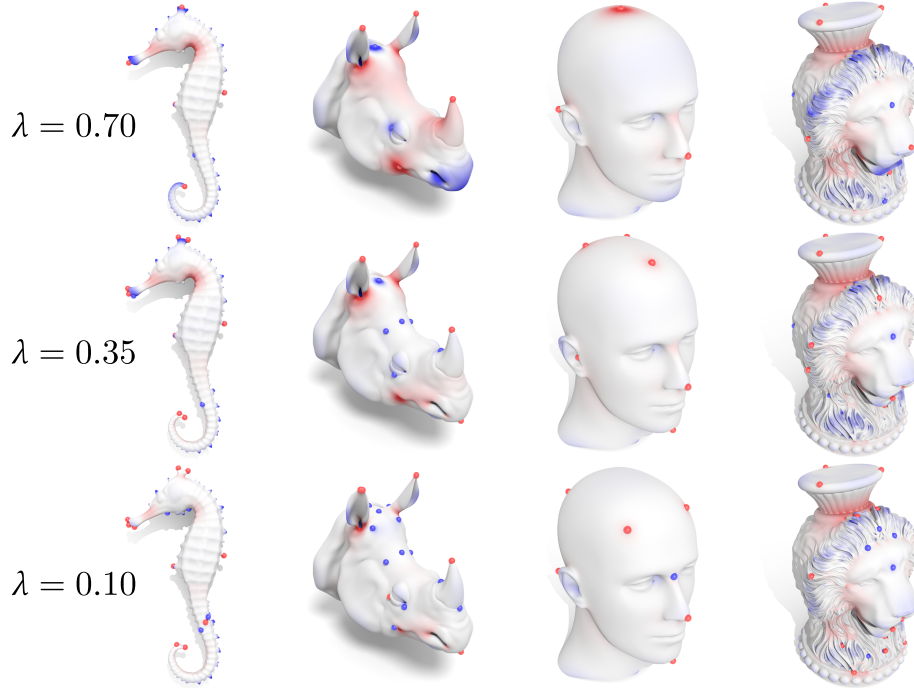


Figure 6.2: The parameter λ has a very consistent behavior across different surfaces, typically producing a similar density of cones (assuming area is normalized to 1).

6.1.2 Comparisons

We here compare the results of our method (**MAD**) to existing cone singularity placement strategies introduced by Ben-Chen et al. [8] (**CPMS**) and Myles & Zorin [45] (**GPIF**). The basic cone placement strategy from Springborn et al. [54] (**CETM**) is similar to **CPMS**, but we omit a comparison since their strategy for picking angles is highly mesh dependent as discussed in Sec. 3.3. **CPMS** takes the target number of cones as input, hence we sometimes show multiple examples. For **GPIF** we do not apply the secondary rounding procedure (which is needed only the special case of integer grid maps), since it would only yield greater area distortion. We also extensively tuned parameters in **GPIF** to achieve the best possible results; we did not try to aim for computational efficiency.

Since one can always reduce area distortion by adding more cones, it is worth thinking about a reasonable way to evaluate the relative “cost” of different cone configurations. One standard approach is to measure the *number* of cones, though on its own this number can be misleading: for instance, as cone angles approach zero they have little real effect on a flattening. Moreover, as suggested by our stability analysis (Sec. 4.4) many small cones placed nearby can have a nearly identical effect to a single large cone of equal total angle. In most examples we therefore report both the number of cones n and the total cone angle $\Phi = \sum_i |\phi_i|$, as well as the resulting L^2 area distortion \mathcal{A} .

In some examples our algorithm places cone singularities in a similar fashion to existing techniques, but typically using fewer cones or smaller total cone angle (Fig. 6.3, *top*). In other examples, we obtain much lower area distortion, or alternatively, comparable distortion with far fewer cones (see for instance Figures 1.2 and 6.4, and 6.3, *bottom left*). In Fig. 6.4, we also see that lines of singularities (as sometimes placed by **GPIF**) do not necessarily yield lower area distortion than simply placing a

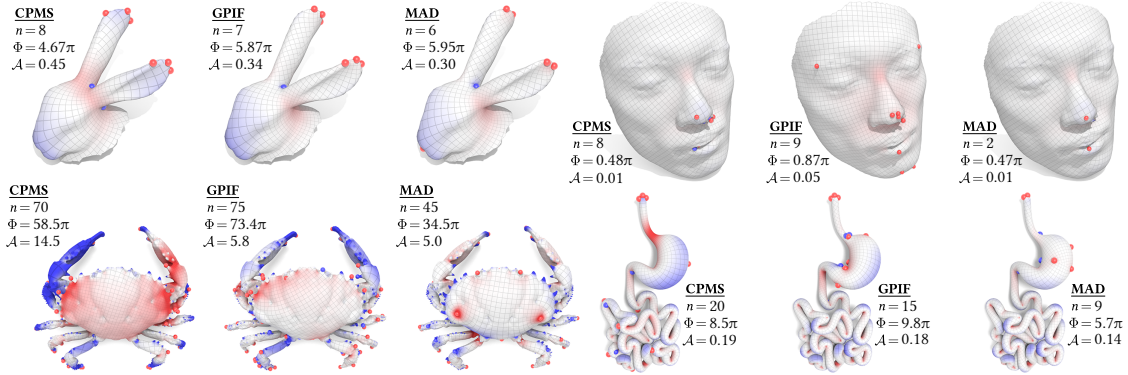


Figure 6.3: More comparisons. On models with simple geometry (*top row*) greedy or region-growing strategies can work quite well, though **MAD** still performs slightly better. On more challenging models such as the crab (*bottom left*) the gap typically widens—note also the high degree of symmetry for **MAD**.




Figure 6.4: Finding *just* the right configuration of cones and angles can sometimes dramatically reduce area distortion. Here, **MAD** almost completely eliminates area distortion using just 8 cones (*far right*). Using the same number of cones in **CPMS** (*far left*) yields far greater distortion; alternatively, one can drive the distortion to similar levels (*center left*) but using far more cones. **GPIF** yields higher distortion than **MAD**, even after placing a whole ring of cones around the top of the head.

few carefully-selected cones. The same example shows that **CPMS** sometimes has a tendency to cluster many cones in the same region, likely due to picking points near the center of a harmonic Green’s function from a prior cone. Overall, we observed similar behavior to these examples across about 50 different meshes of varying geometry, mesh quality, and resolution; in no case did we ever find a configuration with smaller area distortion than **MAD** for equal or smaller total angle Φ .

6.1.3 Robustness

One of the benefits of globally optimal algorithms is that they tend to provide reliable behavior across a larger class of inputs. In Fig. 6.6, we observe that the cones chosen by method are really determined by the geometry of the surface, and are not significantly perturbed by remeshing or common artifacts such as noise, anisotropy, or poor (*e.g.*, non-Delaunay) triangulation. Since we minimize an integral energy, our method is also robust to large outliers, which contribute almost zero area Fig. 6.7. In contrast, **CPMS** will start by placing a cone at *every single outlier*, since they have extremely large scale factors; **GPIF** also puts cones at each of these outliers, since they are (by far) the points of greatest curvature. Finally, Fig. 6.5 demonstrates that **MAD** produces consistent results whether one uses a uniform- or variable-density mesh; in this example, the greedy placement strategy from **CPMS** is confounded by the fact that harmonic Green’s functions will be better resolved—and hence larger—in finer regions, as discussed in Sec. 3.3 (**GPIF** does not suffer from this same artifact).




Figure 6.5: Variable density mesh. *Left:* CPMS places far more singularities in finely tessellated regions, where Green’s functions are better resolved; note that many spheres overlap due to close clustering of cones. *Center:* GPIF also violates symmetry, and achieves lower distortion than MAD only by using about twice as many cones. *Right:* MAD achieves low area distortion using a symmetric arrangement of just a few cones, and with small total cone angle.

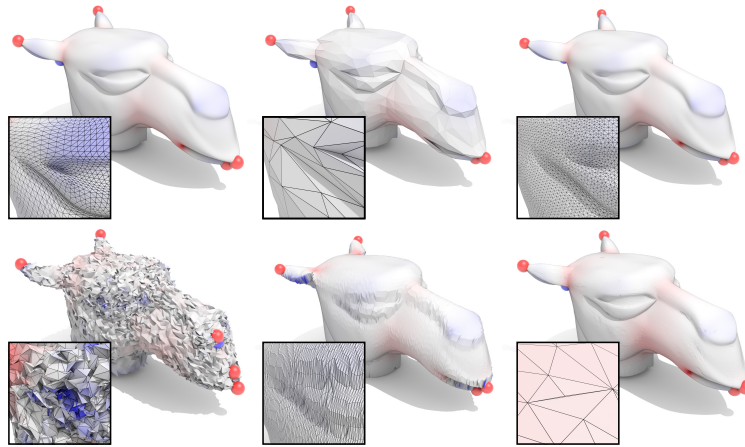


Figure 6.6: Our method produces consistent results on meshes of very different resolution (*top row*) and is also robust to meshing artifacts such as noise (*bottom left*), anisotropy (*bottom center*), and severely non-Delaunay elements (*bottom right*).

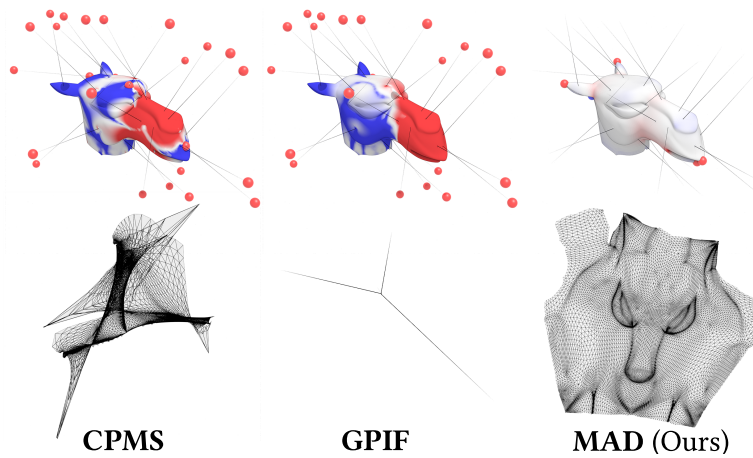


Figure 6.7: Stress test of robustness. Methods that place cones according to peaks in scale factors (*left*) or based on curvature (*center*) may be confounded by outliers; here we consider an extreme case where near-invisible spikes are added to a mesh, leading to cone configurations that are impossible to parameterize (*bottom*). Our method is generally not confounded by outliers or noise, in this case ignoring the spikes and leading to a high-quality parameterization (*right*).

6.2 Extensions

Our basic algorithm can also be extended in a variety of useful ways, as discussed in Sec. 3.5.

6.2.1 Nonuniform Importance




Figure 6.8: *Left:* cones placed according to standard L^2 energy. *Right:* cones placed by re-weighting energy (and regularizer) according to local feature size. Both cone configurations are globally optimal solutions to different problems.

As described in Sec. 3.5, we can augment our method to influence both (i) where cones are placed and (ii) where distortion is measured. More specifically, one can either provide continuous functions $w_{\mathcal{E}}, w_{\mathcal{R}} : V \rightarrow \mathbb{R}_{>0}$ that act as a penalty on distortion and cone placement (*resp.*), or binary functions $U_{\mathcal{E}}, U_{\mathcal{R}} : V \rightarrow \{0, 1\}$ that explicitly excludes regions where distortion is measured and cones are placed (*resp.*). The functions $U_{\mathcal{E}}, U_{\mathcal{R}}$ can be particularly useful for reducing the problem size in cases where there is only a small region of interest—an extreme example is when one wishes to place singularities only along the boundary (see Sec. 6.2.3 for further discussion).

A key example where penalty functions are desirable is on meshes with features across very different scales, such as the fingers and toes on a human body. In this case we first compute a local feature size $r_i := 1/(|\Omega_i|/M_{ii} + \varepsilon)$ at each vertex $i \in V$, where Ω_i is the angle defect (Eqn. 5.2) and M_{ii} is the barycentric dual area (Eqn. 5.1); this value will be small in flat regions and large in highly curved regions. We then set $(w_{\mathcal{E}})_i := r_i$ and $(w_{\mathcal{R}})_i := 1/r_i$, emphasizing the importance of small features, and decreasing the cost of placing cones in those same regions. An example is shown in Fig. 6.8.

An example where excluding a region is natural is when one wants to avoid placing cones in regions that are visible from a particular point of view (Fig. 3.4); here we likewise need only penalize *distortion* in the visible regions. Given a particular viewpoint, we set $(U_{\mathcal{E}})_i$ to 1 and $(U_{\mathcal{R}})_i$ to 0 if and only if vertex i is visible.

Fixed Cone Points with Free Cone Angles In addition to automatically finding the entire configuration, we can optionally allow the user to specify a collection of points $\hat{p}_1, \dots, \hat{p}_m$ that must be included; our method then optimizes the angles of these cones, and also finds the additional cones that best minimize distortion. A critical place where this functionality is needed is finding cone configurations on closed convex surfaces. Consider for instance the unit sphere where there is no reason to place any negative cones—in this scenario, Gauss-Bonnet says that the total cone angle $\Phi := \sum_i |\phi_i|$ will always be 4π . Hence our method will put a cone at *every* vertex i , with cone angle equal to the angle

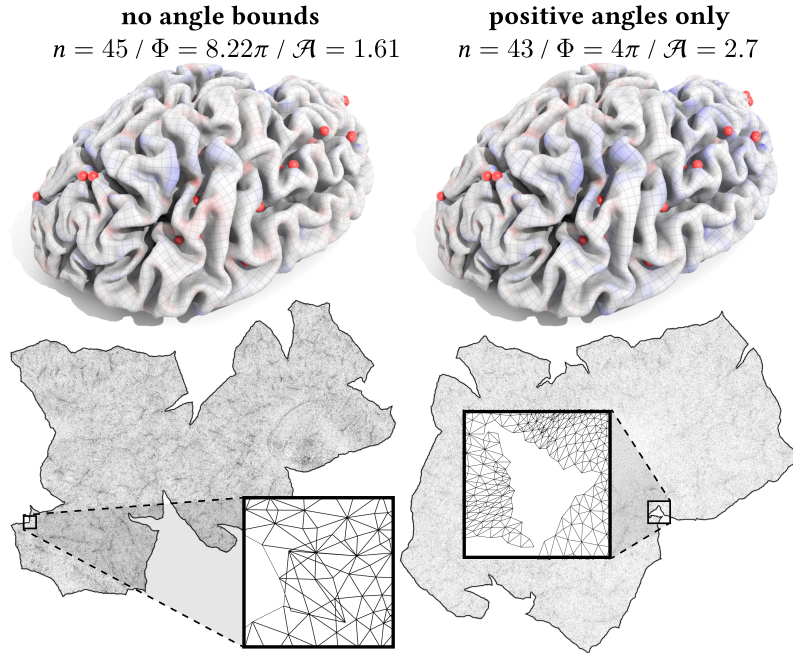


Figure 6.9: *Left:* When finally mapped to the plane a cone flattening of a surface (such as this brain) may have local noninjectivity at negative cones, unless these cones are cut into sufficiently small pieces (see zoom). *Right:* finding an optimal solution with only positive cones avoids this source of local noninjectivity.

defect Ω_i . A simple remedy is to put one “free” cone at an arbitrary vertex $\hat{p} \in V$ (say, the vertex of greatest curvature), which effectively behaves like a small puncture. Our method is now free to find cone configurations where the sum of the cone angle *magnitudes* on the rest of the domain is strictly less than 4π . In practice this strategy is rarely necessary, since most real-world surfaces have both positive and negative curvature.

6.2.2 Bounded Cone Angles

Adding inequality constraints to our optimization (amounting to a simple projection at each iteration) allows us to find optimal configurations with cone angles within a given range. For instance, negative cones can lead to a flattening that is locally noninjective, since the total cone angle is greater than 2π . We can avoid such features by simply requiring that $\phi_i > 0$, helping to improve injectivity. In Fig. 6.9, (*right*) we actually obtain a globally injective flattening, though of course in general one cannot expect global injectivity purely from local injectivity. In one case, we allow a single free cone (as described in Sec. 6.2.1). In another case, we simply optimize over all nonnegative cone configurations with total angle 4π , *without* using including any kind of sparsity-inducing norm—amazingly enough, we still get a sparse solution. *Understanding the sparsity pattern of minimizing positive measures is an important avenue for future research.*


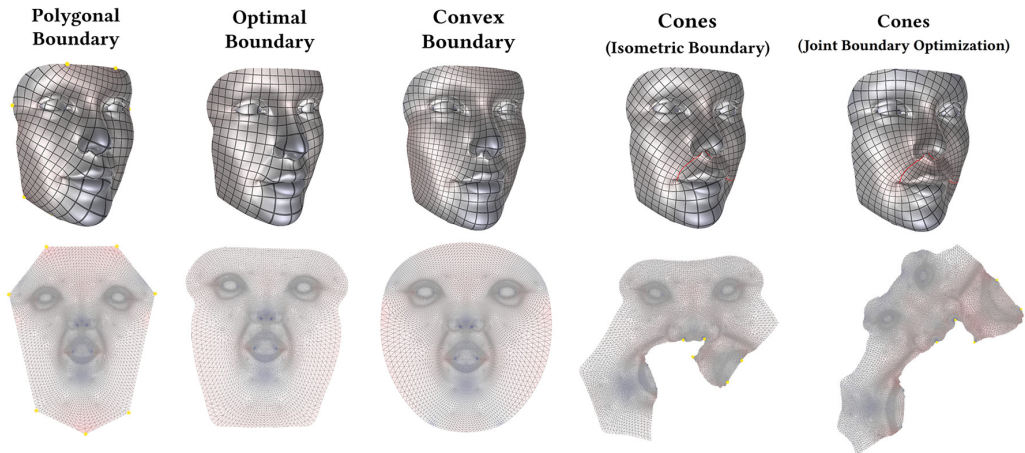


Figure 6.10: *Left:* in the absence of any bounds on cone angles ϕ_i , an optimal configuration for this model is to place eight equal cones corresponding to an octahedron. *Center:* if we now limit angles to the range $\phi \in [0, \pi/2]$, we get a configuration that has eight cones of $\pi/2$ corresponding to a cuboid with unequal lengths. *Right:* the more intuitive configuration with cones at cube corners yields higher area distortion.

Another example where angle bounds are potentially useful is in finding cone configurations for seamless *integer grid maps* [10], where cones must be quantized to integer multiples of $\pi/2$. Although we cannot produce optimal quantized configurations, we find that restricting angles ϕ_i to the range $[-\pi/2, \pi/2]$ often yields a number of $\pm\pi/2$ cones on models that would otherwise have angles outside this range. Fig. 6.10 shows one example where all angles in the optimal configuration do happen to end up being $\pm\pi/2$. Here we observe that the best way to quantize a cone configuration is not always intuitively obvious, indicating there may be significant room for improving existing heuristics found in the meshing literature. Incorporating actual quantization into our framework is therefore an interesting (and challenging) question for future work.

6.2.3 Boundary Treatment

We can get interesting behavior by either encouraging sparsity or bounding angles along the domain boundary, as discussed in Sec. 3.5. For instance, encouraging sparsity of the boundary curvature leads to *polygonal* boundaries; forcing boundary curvature to be positive leads to *convex* boundaries. We can also choose whether to jointly optimize over cone and boundary configurations.




Chapter 7

Discussion and Conclusions

This thesis focused on the problem of determining the optimal configuration of cone singularities for minimizing the area distortion of the resulting conformal cone parameterization. In particular, we formulated the optimal cones problem as a relaxed convex optimization problem with a sparsity inducing regularization. By formulating the optimization problem over the space of finite signed Radon measures we are able to obtain solutions that are supported on Riemannian measure zero sets, like Dirac delta and 1-dimensional Hausdorff measures. By carefully studying a number of prototypical examples from the optimal control community (specifically elliptic and parabolic optimization problems formulated over non-reflexive Banach spaces) we consider an abstract formulation which encapsulates the central features necessary for well-posedness and efficient numerical implementation. Furthermore, the abstract formulation provided a simple framework to verify the correctness of the formalization in the context of Riemannian manifolds—whereas the majority of the work in the optimal control community is formulated over domains in \mathbb{R}^N . From the numerical perspective, we utilized semismooth Newton methods due to their locally superlinear convergence and guarantees regarding mesh-independence. The resulting algorithm is extremely simple to implement using basic tools from numerical linear algebra and mesh processing.

Some of the limitations of our algorithm have already been carefully addressed. For instance, although the continuous problem does not admit exact Delta measures as solutions, we have provided a careful stability analysis (Sec. 4.4) that leads to a practical rounding procedure in the very rare case where cones appear in smaller clusters (Sec. 5.3.5). The numerical experiments indicate that the measure norm regularization *in conjunction* with both the L^2 -norm and the Yamabe equation is a powerful tool for inducing sparsity in optimization—it is important to note that L^1 or measure norms will do very little to promote sparse solutions if the state equation and energy do not prefer such sparse solutions. Another issue is that on surfaces like the unit sphere which have strictly positive Gaussian curvature, the optimal solution to our problem is just the Gaussian curvature measure itself, *i.e.*, a cone at every vertex with angle given by Gauss curvature (see inset, *top*). This measure yields minimal (zero) area distortion, and by Gauss-Bonnet has minimal measure norm. A simple and seemingly effective solution here is to just allow a single “free” cone as described in Sec. 6.2.1 (see inset, *bottom*), though this ‘trick’ is rarely required in practice. There is also some uncertainty in how to pick values of λ , though in practice we find that the same values consistently produce similar results across a wide variety of examples (Fig. 6.2). Perhaps the most interesting question is how to augment our formulation to allow area distortion to be driven below a given user-specified threshold; here some significant new ideas are likely required.



More broadly, despite the importance of the cone flattening problem, very little is known not only about finding optimal solutions, but even about basic questions regarding the behavior of cone flattenings. For instance, there are many outstanding questions about the existence of cone metrics on different topologies or with particular conditions on curvature [22, 23]. One might also wonder about the geometric significance of optimal cone configurations, which might be better understood via connections with optimal transport. From a practical point of view, a major open question is how to find optimal cone configurations where angles are *quantized* (e.g., to integer multiples of $\pi/2$) which are a critical component of structured surface remeshing. Finally, the question of how to optimally drive area distortion below a user-specified bound would enable one to compute high-quality flattenings that are effectively indistinguishable from isometry. Another interesting path for future research is to consider the optimization problem where the L^2 -energy is replaced by the Dirichlet energy. Since the Dirichlet energy of the log-conformal scale factors is infinite in the presence of cone singularities, we should expect that the minimizers are in fact curvature measures supported on curves. These measures supported on curves correspond to *cutting* the surface along these pieces. Preliminary experiments indicate that this is the resulting behavior. The analytical perspectives introduced in this thesis may provide new ways of looking at all of these problems.

Bibliography

- [1] Adams, R. A. & Fournier, J. J. (2003). *Sobolev spaces*, volume 140. Academic press.
- [2] Aigerman, N., Kovalsky, S. Z., & Lipman, Y. (2017). Spherical orbifold tutte embeddings. *ACM Trans. Graph.*, 36(4), 90:1–90:13.
- [3] Aigerman, N. & Lipman, Y. (2015). Orbifold tutte embeddings. *ACM Trans. Graph.*, 34(6), 190:1–190:12.
- [4] Aigerman, N. & Lipman, Y. (2016). Hyperbolic orbifold tutte embeddings. *ACM Trans. Graph.*, 35(6), 217:1–217:14.
- [5] Ambrosio, L., Mainini, E., & Serfaty, S. (2011). Gradient flow of the chapman–rubinstein–schatzman model for signed vortices. In *Annales de l’Institut Henri Poincaré (C) Non Linear Analysis*, volume 28 (pp. 217–246).: Elsevier.
- [6] Aubin, T. (1998). *Some nonlinear problems in Riemannian geometry*. Springer Science & Business Media.
- [7] Bauschke, H. H. & Combettes, P. L. (2017). *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer.
- [8] Ben-Chen, M., Gotsman, C., & Bunin, G. (2008). Conformal flattening by curvature prescription and metric scaling. In *Computer Graphics Forum*, volume 27 (pp. 449–458).: Wiley Online Library.
- [9] Böhm, C. & Ulbrich, M. (2015). A semismooth newton-cg method for constrained parameter identification in seismic tomography. *SIAM J. Scientific Computing*, 37(5).
- [10] Bommès, D., Campen, M., Ebke, H.-C., Alliez, P., & Kobbelt, L. (2013a). Integer-grid maps for reliable quad meshing. *ACM Trans. Graph.*, 32(4), 98:1–98:12.
- [11] Bommès, D., Lévy, B., Pietroni, N., Puppo, E., Silva, C., Tarini, M., & Zorin, D. (2013b). Quad-mesh generation and processing: A survey. Article first published online: 4 MAR 2013, DOI: 10.1111/cgf.12014.
- [12] Brezis, H. (2010). *Functional analysis, Sobolev spaces and partial differential equations*. Springer Science & Business Media.
- [13] Casas, E., Clason, C., & Kunisch, K. (2012). Approximation of elliptic control problems in measure spaces with sparse solutions. *SIAM Journal on Control and Optimization*, 50(4), 1735–1752.
- [14] Casas, E., Clason, C., & Kunisch, K. (2013). Parabolic control problems in measure spaces with sparse solutions. *SIAM J. Control Optim.*, 51(1), 28–63.

- [15] Chao, I., Pinkall, U., Sanan, P., & Schröder, P. (2010). A simple geometric model for elastic deformations. *ACM Trans. Graph.*, 29(4), 38:1–38:6.
- [16] Chen, X., Nashed, Z., & Qi, L. (2000). Smoothing methods and semismooth methods for nondifferentiable operator equations. *SIAM Journal on Numerical Analysis*, 38(4), 1200–1216.
- [17] Clason, C. & Kunisch, K. (2011). A duality-based approach to elliptic control problems in non-reflexive banach spaces. *ESAIM Control Optim. Calc. Var.*, 17(1), 243–266.
- [18] Clason, C. & Kunisch, K. (2012). A measure space approach to optimal source placement. *Computational Optimization and Applications*, 53(1), 155–171.
- [19] Clason, C. & Schiela, A. (2017). Optimal control of elliptic equations with positive measures. *ESAIM: Control, Optimisation and Calculus of Variations*, 23(1), 217–240.
- [20] D’Aprile, T., De Marchis, F., & Ianni, I. (2016). Prescribed gauss curvature problem on singular surfaces. *arXiv preprint arXiv:1612.03657*.
- [21] de Goes, F., Desbrun, M., & Tong, Y. (2015). Vector field processing on triangle meshes. In *SIGGRAPH Asia 2015 Courses, SA ’15* (pp. 17:1–17:48). New York, NY, USA: ACM.
- [22] De Marchis, F. & López-Soriano, R. (2016). Existence and non existence results for the singular nirenberg problem. *Calculus of Variations and Partial Differential Equations*, 55(2), 36.
- [23] Del Pino, M. & Román, C. (2015). Large conformal metrics with prescribed sign-changing gauss curvature. *Calculus of Variations and Partial Differential Equations*, 54(1), 763–789.
- [24] Erickson, J. & Har-Peled, S. (2004). Optimally cutting a surface into a disk. *Discrete & Computational Geometry*, 31(1), 37–59.
- [25] Folland, G. B. (1984). *Real analysis*.
- [26] Fonseca, I. & Leoni, G. (2007). *Modern Methods in the Calculus of Variations: L^p Spaces*. Springer Science & Business Media.
- [27] Gilbarg, D. & Trudinger, N. S. (2001). *Elliptic Partial Differential Equations of Second Order*. Springer.
- [28] Günther, A. & Tber, M. H. (2016). A goal-oriented adaptive moreau-yosida algorithm for control- and state-constrained elliptic control problems. *Advances in Applied Mathematics and Mechanics*, 8(3), 426–448.
- [29] Helms, L. L. V. (2009). *Potential theory*. Springer.
- [30] Hencky, H. (1928). Über die form des elastizitätsgesetzes bei ideal elastischen stoffen. *Zeitschrift für technische Physik*, 9, 215–220.
- [31] Hintermüller, M. (2010). Semismooth newton methods and applications. *Department of Mathematics, Humboldt-University of Berlin*.

- [32] Hintermüller, M., Ito, K., & Kunisch, K. (2002). The primal-dual active set strategy as a semi-smooth newton method. *SIAM J. on Optimization*, 13(3), 865–888.
- [33] Hintermüller, M. & Kunisch, K. (2006). Feasible and noninterior path-following in constrained minimization with low multiplier regularity. *SIAM Journal on Control and Optimization*, 45(4), 1198–1221.
- [34] Hinze, M. (2005). A variational discretization concept in control constrained optimization: The linear-quadratic case. *Comp. Opt. and Appl.*, 30(1), 45–61.
- [35] Ito, K. & Kunisch, K. (2003a). Semi-smooth newton methods for state-constrained optimal control problems. *Systems & Control Letters*, 50(3), 221–228.
- [36] Ito, K. & Kunisch, K. (2003b). Semi-smooth newton methods for variational inequalities of the first kind. *ESAIM: Mathematical Modelling and Numerical Analysis*, 37(1), 41–62.
- [37] Ito, K. & Kunisch, K. (2008). Semi-smooth newton methods for the signorini problem. *Applications of Mathematics*, 53(5), 455–468.
- [38] Jost, J. (2017). *Riemannian Geometry and Geometric Analysis*. Universitext. Springer.
- [39] Kharevych, L., Springborn, B., & Schröder, P. (2006). Discrete conformal mappings via circle patterns. *ACM Trans. Graph.*, 25(2), 412–438.
- [40] Knöppel, F., Crane, K., Pinkall, U., & Schröder, P. (2013). Globally optimal direction fields. *ACM Trans. Graph.*, 32(4).
- [41] Konakovic, M., Crane, K., Deng, B., Bouaziz, S., Piker, D., & Pauly, M. (2016). Beyond developable: Computational design and fabrication with auxetic materials. *ACM Trans. Graph.*, 35(4).
- [42] MacNeal, R. (1949). *The Solution of Partial Differential Equations by Means of Electrical Networks*. PhD thesis, California Institute of Technology.
- [43] Mainini, E. (2012). A description of transport cost for signed measures. *Journal of Mathematical Sciences*, 181, 837–855.
- [44] Maron, H., Galun, M., Aigerman, N., Trope, M., Dym, N., Yumer, E., Kim, V. G., & Lipman, Y. (2017). Convolutional neural networks on surfaces via seamless toric covers. *ACM Trans. Graph.*, 36(4), 71:1–71:10.
- [45] Myles, A. & Zorin, D. (2012). Global parametrization by incremental flattening. *ACM Transactions on Graphics (TOG)*, 31(4), 109.
- [46] Myles, A. & Zorin, D. (2013). Controlled-distortion constrained global parametrization. *ACM Transactions on Graphics (TOG)*, 32(4), 105.
- [47] Rockafellar, R. T. (1997). *Convex analysis*. princeton landmarks in mathematics.
- [48] Saff, E. & Totik, V. (1997). *Logarithmic Potentials with External Fields*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg.

- [49] Sawhney, R. & Crane, K. (2017). Boundary first flattening. *ACM Trans. Graph.*, 37(1), 5:1–5:14.
- [50] Schmidt, R. & Singh, K. (2010). Meshmixer: An interface for rapid mesh composition. In *ACM SIGGRAPH 2010 Talks*, SIGGRAPH '10 (pp. 6:1–6:1). New York, NY, USA: ACM.
- [51] Sheffer, A., Praun, E., & Rose, K. (2006). Mesh parameterization methods and their applications. *Found. Trends. Comput. Graph. Vis.*, 2(2), 105–171.
- [52] Song, X., Chen, B., & Yu, B. (2017). An efficient duality-based approach for pde-constrained sparse optimization. *Computational Optimization and Applications*, (pp. 1–40).
- [53] Sorkine, O. & Alexa, M. (2007). As-rigid-as-possible surface modeling. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*, SGP '07 (pp. 109–116). Aire-la-Ville, Switzerland, Switzerland: Eurographics Association.
- [54] Springborn, B., Schröder, P., & Pinkall, U. (2008). Conformal equivalence of triangle meshes. In *ACM Transactions on Graphics (TOG)*, volume 27 (pp.77):. ACM.
- [55] Stampacchia, G. (1965). Le problème de dirichlet pour les équations elliptiques du second ordre à coefficients discontinus. *Ann. Inst. Fourier (Grenoble)*, 15(1), 189–258.
- [56] Thurston, W. (2002). The Geometry and Topology of 3-Manifolds.
- [57] Thurston, W. P. (1998). Shapes of polyhedra and triangulations of the sphere. *ArXiv Mathematics e-prints*.
- [58] Tröltzsch, F. & Yousept, I. (2009). A regularization method for the numerical solution of elliptic boundary control problems with pointwise state constraints. *Computational Optimization and Applications*, 42(1), 43–66.
- [59] Troyanov, M. (1989). Metrics of constant curvature on a sphere with two conical singularities. *Lect. Notes Math*, 1410, 296–308.
- [60] Troyanov, M. (1991). Prescribing curvature on compact surfaces with conical singularities. *Transactions of the American Mathematical Society*, 324(2), 793–821.
- [61] Tsui, A., Fenton, D., Vuong, P., Hass, J., Koehl, P., Amenta, N., Coeurjolly, D., DeCarli, C., & Carmichael, O. (2013). Globally Optimal Cortical Surface Matching With Exact Landmark Correspondence. In *Information Processing in Medical Imaging*, LNCS (pp. 487–498):. Springer-Verlag.
- [62] Ulbrich, M. (2002a). *Nonsmooth Newton-like Methods for Variational Inequalities and Constrained Optimization Problems in Function Spaces*, Habilitationsschrift. Germany: Fakultät für Mathematik, Technische Universität München.
- [63] Ulbrich, M. (2002b). Semismooth newton methods for operator equations in function spaces. *SIAM J. Optim.*, 13(3), 805–841.

- [64] Ulbrich, M. (2011). *Semismooth Newton Methods for Variational Inequalities and Constrained Optimization Problems in Function Spaces*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics.
- [65] Vaxman, A., Campen, M., Diamanti, O., Panozzo, D., Bommers, D., Hildebrandt, K., & Ben-Chen, M. (2016). Directional Field Synthesis, Design, and Processing. *Computer Graphics Forum*.
- [66] Villani, C. (2009). *Optimal transport: old and new*, volume 338. Springer Science & Business Media.
- [67] Vintescu, A., Dupont, F., Lavoué, G., Memari, P., & Tierny, J. (2017a). Least Squares Affine Transitions for Global Parameterization. In *WSCG 2017 5th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision 2017*.
- [68] Vintescu, A.-M., Dupont, F., Lavoué, G., Memari, P., & Tierny, J. (2017b). Conformal Factor Persistence for Fast Hierarchical Cone Extraction. In A. Peytavie & C. Bosch (Eds.), *EG 2017 - Short Papers*: The Eurographics Association.