

Optimal Cone Singularities for Conformal Flattening

YOUSUF SOLIMAN, DEJAN SLEPČEV, and KEENAN CRANE, Carnegie Mellon University

Angle-preserving or *conformal* surface parameterization has proven to be a powerful tool across applications ranging from geometry processing, to digital manufacturing, to machine learning, yet conformal maps can still suffer from severe area distortion. *Cone singularities* provide a way to mitigate this distortion, but finding the best configuration of cones is notoriously difficult. This paper develops a strategy that is globally optimal in the sense that it minimizes total area distortion among all possible cone configurations (number, placement, and size) that have no more than a fixed total cone angle. A key insight is that, for the purpose of optimization, one should not work directly with curvature measures (which naturally represent cone configurations), but can instead apply *Fenchel-Rockafellar duality* to obtain a formulation involving only ordinary functions. The result is a convex optimization problem, which can be solved via a sequence of sparse linear systems easily built from the usual cotangent Laplacian. The method supports user-defined notions of importance, constraints on cone angles (e.g., positive, or within a given range), and sophisticated boundary conditions (e.g., convex, or polygonal). We compare our approach to previous techniques on a variety of challenging models, often achieving dramatically lower distortion, and demonstrating that global optimality leads to extreme robustness in the presence of noise or poor discretization.

CCS Concepts: • **Computing methodologies** → **Mesh geometry models**; • **Mathematics of computing** → *Continuous optimization*;

Additional Key Words and Phrases: discrete differential geometry, geometry processing, conformal mapping, mesh parameterization

ACM Reference Format:

Yousuf Soliman, Dejan Slepčev, and Keenan Crane. 2018. Optimal Cone Singularities for Conformal Flattening. *ACM Trans. Graph.* 37, 4, Article 105 (August 2018), 24 pages. <https://doi.org/10.1145/3197517.3201367>

1 INTRODUCTION

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 [Konakovic et al. 2016; Maron et al. 2017]. 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




Fig. 1. *Left*: Conformal flattenings have no angle distortion, but can severely distort area. *Right*: When placed optimally, even a small number of cone singularities (here just nine) can almost completely eliminate area distortion.

is represented by only a tiny region in the parameter domain. The basic idea behind *cone flattening* [Kharevych et al. 2006] 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. 2). 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, determined 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 the problem well-posed, one can fix the number of cones, or alternatively (as we will do), fix the total magnitude $\Phi := \sum_i |\phi_i|$ of all cone angles ϕ_i .



Fig. 2. A conformal cone parameterization is equivalent to flattening a smooth surface, like the sphere, over a polyhedron, which can then 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.)

Authors' address: Yousuf Soliman; Dejan Slepčev; Keenan Crane, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2018/8-ART105 \$15.00 <https://doi.org/10.1145/3197517.3201367>

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. Our method ensures that total area distortion is always globally minimized, providing substantial robustness and, in practice, significantly lower distortion on many real examples. Ultimately, 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 easily accelerated with a simple multi-resolution scheme,
- provides user control over regions where cones can be placed, as well as regions where distortion should be penalized,
- 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 $[-\pi/2, \pi/2]$).

Since we achieve *minimal area distortion*, we will use the acronym **MAD** throughout to refer to our method. On the whole, flattenings produced this way are extremely close to isometric, making them broadly useful for a variety of practical applications.

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 paper and in the development of future work. In particular, we look at the practical importance of choosing a principled measure of area distortion (Sec. 4.1.2), we analyze stability of cone flattenings with respect to perturbations of the cones, *i.e.*, how much will distortion change if singularities are “merged”; and we consider the approximability of smooth metrics by polyhedral metrics from an analytical point of view, *i.e.*, when can a given metric be arbitrarily well-approximated by cones (App. C). 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 space H^{-1} and hence cannot exactly describe cones, which correspond to delta measures. In practice we can (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. 9).

Though our method takes some work to develop, the final algorithm can be easily implemented using standard tools from geometry processing; a practical description is presented in Sec. 5.3. To derive this algorithm, we start in the smooth setting (Sec. 4) where we formulate the task as a PDE-constrained optimization problem involving the *Yamabe equation* from conformal geometry. Since this problem is nonconvex, we formulate a convex relaxation over a larger space of curvature measures. This problem has the right minimizers, but poses some numerical challenges; we therefore formulate its *Fenchel-Rockafellar dual*, which is easily discretized via ordinary finite elements (Sec. 5). We numerically validate our findings and explore comparisons with prior work in Sec. 7.

REMARK. *The appendices, which include detailed computations, derivations, and proofs, can be found in the supplemental material. An extended version of this work can be found in the MS thesis of the first author [Soliman 2018].*

2 RELATED WORK

A variety of problems in digital geometry processing seek ideal locations for certain singular features. For instance, in the design of tangent vector fields or rotationally symmetric direction fields, judicious placement of singularities can significantly impact the global regularity of the field [Vaxman et al. 2016]. In this setting one can find optimal singularities by simply solving a sparse eigenvalue problem [Knöppel et al. 2013]. However, any connection to the problem we consider here is fairly indirect: although such fields are often used in parameterization algorithms [Bommes et al. 2013b], singularities that yield highly regular fields do not immediately guarantee a good flattening (Fig. 17). There is also an extensive literature on *cutting* surfaces into pieces that can be flattened with low area distortion [Sheffer et al. 2006]; an important distinction between general cuts and those arising from a cone flattening is that the latter is automatically *seamless* away from a small set of isolated points—helping to avoid common filtering artifacts, and facilitating physical fabrication. To date there is no general cutting method that guarantees globally minimal area distortion (though local optimality has recently been studied [Sharp and Crane 2018]); in fact even the simpler problem of finding the *shortest* way of cutting a surface into a disk is NP-hard [Erickson and Har-Peled 2004].

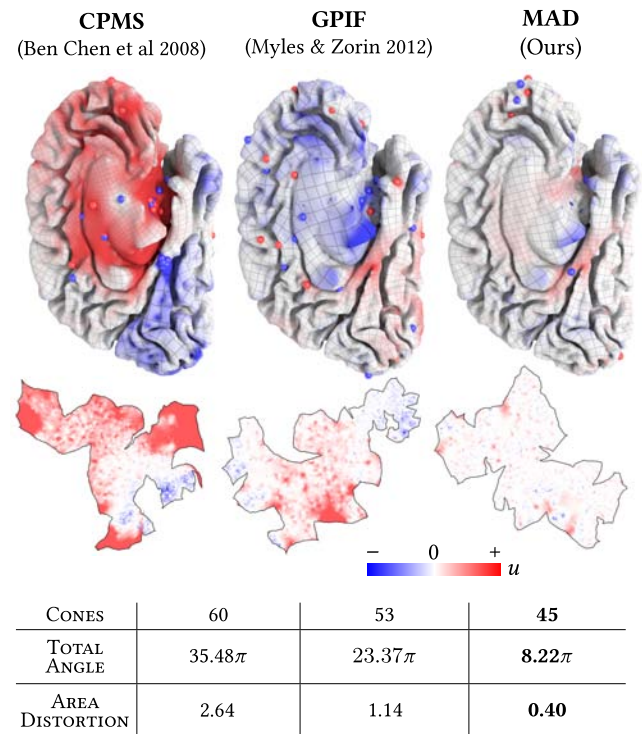


Fig. 3. Even with fewer cones and much smaller total cone angle, our cone placement strategy (**MAD**) yields far lower area distortion than previous methods. This effect is especially apparent on shapes like the brain, which do not have obvious peaks of curvature.

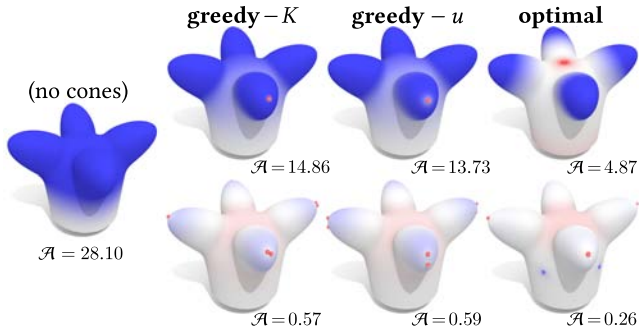


Fig. 4. Intuitively, it might seem that the best place to put cones is in regions of high curvature K or large scale distortion u . However, the optimal location may actually be a point that is nearly flat—helping to explain the suboptimal behavior of greedy algorithms. Here we place either one cone (*top*) or eight cones (*bottom*); notice that even for a larger number of cones, the optimal configuration continues to include cones in flat regions.

Conformal Cone Singularity Placement. In this paper, we focus specifically on cone singularities in the context of conformal flattening; here a variety of strategies have been proposed. Kharevych et al. [2006] initially investigated cone flattenings by manually drawing layouts and adjusting cone angles to reduce distortion. Springborn et al. [2008] 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 at the point of greatest area distortion. In subsequent iterations, cone points are treated as punctures in the domain, leading to cone angles that are automatically determined by the flattening process. As we will discuss in Sec. 4.1.2, however, this approach is mesh dependent since the Dirichlet energy of the log conformal factor blows up in the presence of cones. In a parallel development, Ben-Chen et al. [2008] devise a flattening algorithm (**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. [2017b] using hierarchical persistence. Later, Myles and Zorin [2012] developed a method for seamless global parameterization (**GPIF**), where cones are determined by incrementally flattening the surface starting with the regions of 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. 4. Another example of how curvature-based approaches may lead to suboptimal solutions is shown in Fig. 3, where curvature is not distributed around any obvious peaks. 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. The nonlocal nature of this problem arises from the fact that these functions do not have compact support.

Cone Metrics, Orbifolds, and Liouville Equations. Cone singularities can be understood from several different points of view. Thurston [2002, Chapter 13] studied a geometric picture of manifolds with an *orbifold structure*, i.e., each point must locally look Euclidean or like the quotient of a Euclidean space under a discrete group action. Very recently this orbifold perspective has become quite fruitful in algorithms, leading to methods for computing canonical mappings between surfaces with landmarks [Tsui et al. 2013], and parameterization algorithms with guarantees on global injectivity [Aigerman and Lipman 2015, 2016; Aigerman et al. 2017]. A very different perspective centers around the analytical viewpoint 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 early foundations for studying this equation in the context of singular cone metrics [Troyanov 1989, 1991], which continues to be investigated [Del Pino and Román 2015; De Marchis and López-Soriano 2016; D’Aprile et al. 2016]. This intrinsic, analytic point of view serves as a starting point for many recent algorithms including **CPMS**, **CETM**, and **GPIF**, as well as a recent method for conformal flattening [Sawhney and Crane 2017]; it also plays a fundamental role in the method we develop here.

Convex Optimization and Semismooth Operator Equations. Carefully formulating the cone placement problem using the tools of analysis allows us to take advantage of highly effective methods for *semismooth operator equations* recently developed by the optimal control community. A key insight of these methods is that one should not directly discretize the original problem, but rather formulate optimality conditions in the continuous setting, *then* discretize these conditions [Ito and Kunisch 2003b; Günther and Tber 2016]. 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 discrete, finite dimensional spaces. In our case, a primal/dual formulation yields optimality conditions that are nicely solved via *semismooth Newton methods* [Chen et al. 2000; Ulbrich 2002, 2011]. In particular, our approach is similar to the approach of Hinze [2005], 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 quite 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 [Hintermüller et al. 2002; Ito and Kunisch 2003a]. 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 [Clason and Kunisch 2012; Casas et al. 2012; Clason and Schiela 2017]. Our work builds on this literature and shows how similar formulations can be applied to problems in geometry processing and computer graphics.

3 BACKGROUND

We briefly review concepts from conformal geometry and measure theory that will be needed to develop our formulation (for a more comprehensive introduction, see [Bauschke and Combettes 2017]). The reason for discussing these particular concepts in detail is that, in the smooth setting, our problem must be formulated in terms of *measures*—but when it comes to computation, we really only know how to discretize *functions*. Sec. 3.2 explains how the space of measures can be viewed as the *dual* of a space of functions, and Sec. 3.4 subsequently explains how an optimization problem in one space can be turned into a problem involving the dual space—in particular, allowing us to transition from measures to functions.

Throughout, the domain for our problem is a smooth surface M with Riemannian metric g , and corresponding area element dA . We use the notation $\bar{\mathbb{R}} := \mathbb{R} \cup \{+\infty\}$.

3.1 Conformal Flattening

A Riemannian metric \tilde{g} is *conformally equivalent* to g if it can be expressed as 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 a 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*:

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

where Δ is the Laplace-Beltrami operator on (M, g) , and K, \tilde{K} are the initial and target Gaussian curvature, *resp.* Along the boundary ∂M , the change in geodesic curvature κ is described by the *Cherrier boundary conditions*:

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

where \mathbf{n} is the normal to the domain boundary. (A derivation of these equations can be found in [Aubin 1998].)

A *polyhedral cone metric* is a Riemannian metric that locally looks either flat, or like the tip of a cone—the chief example being the metric of any Euclidean polyhedron, as depicted in Fig. 2. We will use $p_1, \dots, p_m \in M$ to denote the cone points (corresponding to the vertices of a polyhedron), and $\theta_1, \dots, \theta_m$ to refer to the angle deficit obtained if the cone is cut open and isometrically flattened—*e.g.*, $\theta = 0$ at flat points (see [Trojanov 1989, Definition 1] for a more formal description).

When conformally mapping a smooth surface to a polyhedral cone metric, the Yamabe equation becomes

$$\begin{cases} \Delta u = K - \sum_{i=1}^m \theta_i \delta_{p_i} & \text{in } M, \\ \frac{\partial u}{\partial \mathbf{n}} = \kappa_g - e^u \kappa_{\tilde{g}} & \text{on } \partial M, \end{cases} \quad (1)$$

where δ_p denotes the Dirac delta measure at p . Eqn. 1 serves as the key starting point in developing a convex global optimization problem for cone placement.

3.2 Measure Spaces

A key feature of our approach is that the underlying optimization problem is framed in terms of *measures*, rather than ordinary functions. Loosely speaking, a measure assigns a “size” to subsets of the domain in a natural way—for example, if U is a region of the plane




Fig. 5. Near a cone singularity, the solution u to the Yamabe equation looks like a harmonic Green’s function (*far right*). Methods that pick cones based on, *e.g.*, peaks of the scale factor have trouble in the discrete setting, since mesh resolution completely changes the behavior of this function (*top row*).

\mathbb{R}^2 , then $\mu(U) = \int_U dx dy$ is the usual area of U . In general, any measure μ defines a notion of integration: for example, the expression $\int_M d\mu$ gives the total area with respect to μ ; $\int_M f d\mu$ integrates values of f over M , weighted by μ . Note that in general μ can also assign negative values—more formally, we will work with finite signed Radon measures on M ; we will use $\mathcal{M}(M)$ to denote the set of all such measures on M , and use $\int_M \cdot d\mu$ to denote the Lebesgue integral with respect to μ .

Linear Duality. Let $C(M)$ denote the space of all continuous functions $\varphi : M \rightarrow \mathbb{R}$; a *functional* on $C(M)$ is a map $L(\varphi)$ that assigns a real value to any such function. A simple example is integration of φ over the whole domain, *i.e.*, $L(\varphi) = \int_M \varphi dA$. The *Riesz representation theorem* says that the space of measures is equivalent to the space of (*continuous*) *linear functionals* L on $C(M)$. In other words, for any such functional, there exists a unique measure $\mu \in \mathcal{M}(M)$ such that

$$L(\varphi) = \int_M \varphi d\mu, \quad \text{for all } \varphi \in C(M)$$

(and vice versa). For instance, for standard integration with respect to surface area, we just have $d\mu = dA$. More generally, for any function f , we have a measure $d\mu = f dA$ corresponding to the functional $L(\varphi) = \int_M \varphi f dA$, *i.e.*, an integral weighted by the function f . A more interesting example is the (*Dirac*) *delta measure* δ_p associated with the linear functional $L(\varphi) = \varphi(p)$, *i.e.*, the functional that simply yields the value of the function at the point p —importantly, this measure cannot be represented as $f dA$ for any function f . This identification between measures and linear functionals turns out to be essential for formulating dual optimization problems, as discussed in Sec. 3.4.

Measure Norm. We can quantify the overall magnitude of a measure μ via its *measure norm*

$$\|\mu\|_{\mathcal{M}} := \sup_{\varphi \in C(M)} \left\{ \int_M \varphi d\mu : |\varphi(x)| \leq 1 \text{ for all } x \in M \right\}.$$

For instance, if μ is a positive measure, then $\|\mu\|_{\mathcal{M}}$ is just the area of M with respect to μ . If μ is a weighted sum of delta measures $\mu = \sum_i \theta_i \delta_{p_i}$ (as in our cone placement problem), then the measure norm is just the ℓ^1 norm $\|\mu\|_{\mathcal{M}} = \sum_i |\theta_i|$; likewise, when $d\mu = f dA$, the norm $\|\mu\|_{\mathcal{M}}$ amounts to the L^1 norm of f , *i.e.*, $\int_M |f| dA$. These examples hint at the fact that $\|\mu\|_{\mathcal{M}}$ provides a way to encourage sparsity in optimization problems involving more general measures (see Sec. 4.2 for further discussion).

3.3 Linear and Convex Analysis

Some elementary ideas from analysis will be needed to formulate our dual optimization problem. First, any normed vector space X has a *dual vector space* X^* consisting of continuous linear maps L from X to \mathbb{R} . For any linear map $\Lambda : X \rightarrow Y$ between vector spaces X and Y , its (*Banach space*) *adjoint* is a linear map $\Lambda^* : Y^* \rightarrow X^*$ such that for all vectors $x \in X$ and linear functionals $L : Y \rightarrow \mathbb{R}$, $(\Lambda^*L)(x) = L(\Lambda x)$. In finite dimensions, for instance, the adjoint is represented by the matrix transpose. Finally, the *convex conjugate* of any function $F : X \rightarrow \overline{\mathbb{R}}$ is the convex 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 (see App. A.3).

3.4 Convex Duality

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 [Soliman 2018, Chapter 4.3]). We instead use the more specialized technique of *Fenchel-Rockafellar duality*, which for problems of the kind considered in this paper easily yields an explicit characterization of minimizers. An excellent reference on Fenchel-Rockafellar duality is Brezis [2010].

In particular, suppose we want to solve the problem

$$\inf_{x \in X} F(x) + G(\Lambda x), \quad (2)$$

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 convex problem

$$\max_{y^* \in Y^*} F^*(\Lambda^* y^*) - G^*(y^*), \quad (3)$$

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

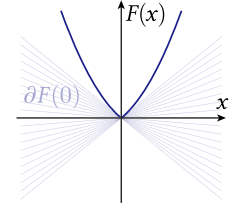
Optimality Conditions. As with Lagrange duality, optimal points can be nicely characterized in terms of both primal and dual variables. For a differentiable objective, these conditions would simply involve derivatives of F and G . If they are not differentiable—as will be the case in our problem (Sec. 4.2)—we can instead 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. More precisely, for any convex function F from a normed vector space X to \mathbb{R} , the subdifferential ∂F 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 the set of slopes $\partial F(0) = [-1, 1]$.

Using the subdifferential, we can express the optimality conditions as

$$\begin{cases} \Lambda^* \bar{y}^* \in \partial F(\bar{x}), \\ \bar{y}^* \in \partial G(\Lambda \bar{x}). \end{cases} \quad (4)$$



For example, when Λ is the identity and both F and G are differentiable, these conditions amount to saying that $\nabla F(\bar{x}) = \nabla G(\bar{x})$, similar to 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 yields numerical behavior far superior to solving the primal or dual problem directly (Sec. 4.2). 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 appears to preserve essential structure from the continuous setting (namely, relationships between primal and dual variables). Similar observations have recently been made in the context of optimal control [Roland and Karl 2010].

4 SMOOTH FORMULATION

Our basic goal is to find a target cone metric that yields least area distortion under a conformal map. For any given configuration of cones, this distortion is determined by the singular Yamabe equation (Eqn. 1), leading to a PDE-constrained optimization problem (Sec. 4.1). In order to solve this problem, we then consider a convex relaxation (Sec. 4.2) which yields essentially optimal solutions to the original problem, as discussed in Sec. 4.5. Finally, Sec. 4.6 introduces regularization that helps with numerical stability and efficiency.

4.1 Basic Problem

Our basic problem is to find a collection of cone points $p_i \in M$ and corresponding cone angles $\theta_i \in \mathbb{R}$ that minimize area distortion under the conformal rescaling $\tilde{g} = e^{2u}g$ that takes us from the original metric g to the polyhedral cone metric \tilde{g} . Since the scale factor u is determined by the singular Yamabe equation (Eqn. 1), we can formulate this problem as

$$\begin{array}{ll} \text{minimize} & E(u) \\ p_i \in M, \theta_i \in \mathbb{R} & \\ \text{subject to} & \Delta u = K \quad \sum_i \theta_i \delta_{p_i} \quad \text{in } M \\ & u = 0 \quad \text{on } \partial M, \end{array} \quad (5)$$

where the energy E quantifies the overall area distortion—we will discuss particular choices of energy in Sec. 4.1.2. There are two issues with this problem as stated. First, although it is convex in the angles θ_i , it is not convex in the positions p_i . Second, it is ill-posed in the sense that one can always reduce area distortion further by adding more cones. Both of these issues will be addressed via a relaxation introduced in Sec. 4.2.

4.1.1 Local Picture. A simpler, local picture helps with both intuition and analysis. Here we assume (M, g) is a topological disk, which is initially flattened to the plane via conformal scale factors u_0 . We then seek subsequent scaling by factors u that yield a low-distortion cone metric

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

In this setting, Problem 5 becomes

$$\begin{aligned} & \text{minimize} && E(u + u_0) \\ & \text{subject to} && \Delta_{\mathbb{R}^2} u = \sum_i \theta_i \delta_{p_i} \quad \text{in } M \\ & && u = 0 \quad \text{on } \partial M. \end{aligned} \quad (6)$$

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

This local point of view provides a different perspective on the cone placement problem: the function u_0 describes the scale distortion of the initial flattening, and we seek the best approximation of this function by a weighted sum of *harmonic Green's functions*, i.e., solutions to $\Delta_{\mathbb{R}^2} u = \delta_p$ (as pictured in Fig. 5). Simply picking extrema of Gaussian curvature or peaks in the initial scale factor (as in greedy strategies) will not in general yield optimal cone placement, since the long tails of harmonic Green's functions can substantially influence the result (Fig. 4, *left*). Conversely, cones carefully arranged in flat regions can conspire to reduce distortion in regions of greater curvature (Fig. 4, *right*). Of course, for the problem of best approximation to be meaningful we must first answer the question of what it means for an approximation to be “best.”

4.1.2 Measures of Area Distortion. How do we measure the area distortion of a conformal flattening? Springborn et al. [2008] 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 log conformal factors:

$$E_D(u) := \frac{1}{2} \int_M |\nabla u|^2 dA.$$

When this energy is small, it indicates that scaling is near-constant, i.e., low area distortion up to uniform scaling. However, this energy is not meaningful in the context of cone flattening since scale factors blow up logarithmically near a singularity. 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 is highly mesh-dependent, since scale factors due to cones will be better resolved—and hence much larger—in densely sampled regions (see Fig. 5). Any method that aims to minimize E_D will therefore prefer to place cones in coarsely sampled regions, even if they are not geometrically relevant. Likewise, the cone angles chosen by CETM go to zero under refinement, since large cones become increasingly expensive (Fig. 6).

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

$$E(u) := \int_M u^2 dA.$$

This energy is finite for any solution of the singular Yamabe equation, and hence converges to a finite value under mesh refinement. In mechanics, E is known as the *true strain* [Hencky 1928]; as noted

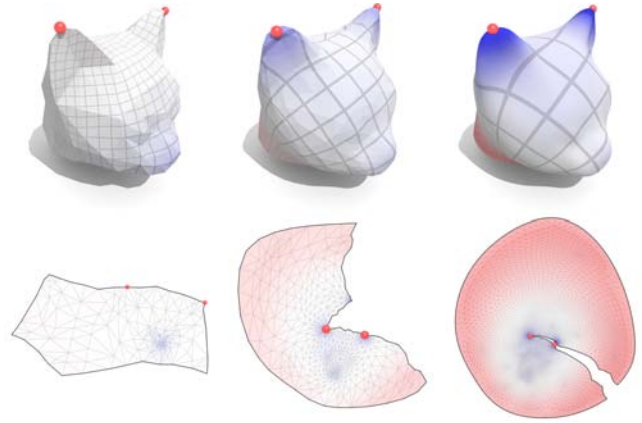


Fig. 6. Since the Dirichlet energy of a harmonic Green's function blows up under mesh refinement, cone angles obtained in CETM by setting $u = 0$ at the point of maximum distortion will tend to zero for fine meshes.

by Myles and Zorin [2012], it is also a second-order approximation of a nonlinear elastic energy [Chao et al. 2010], known in computer graphics as the *as rigid as possible* energy [Sorkine and Alexa 2007].

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. A nice alternative we consider in this paper is re-weighting the L^2 energy by a local feature size (Sec. 8.1); Fig. 7 shows one example.

4.2 Relaxation

We now introduce a relaxation that addresses the two problems mentioned in Sec. 4.1, namely, nonconvexity and ill-posedness. Recall in particular that the cone placement problem (Problem 5) has no solution: we can make the area distortion arbitrarily small by simply introducing more and more cones. In the discrete case, for instance, one could simply put a cone at every vertex. Limiting the *number* of cones yields (in the discrete setting) a combinatorial

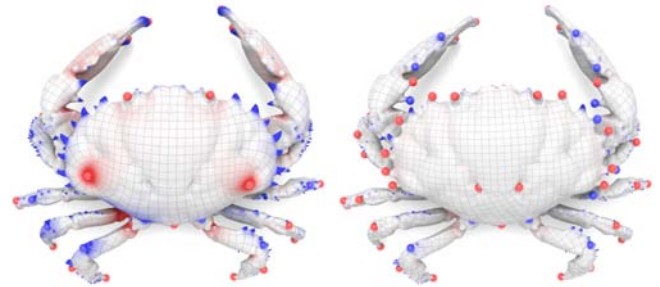


Fig. 7. *Left*: cones placed according to standard L^2 energy. *Right*: cones placed by re-weighting the L^2 energy (and regularizer) according to local feature size. Both cone configurations are globally optimal solutions to different problems.

optimization problem which appears difficult to solve optimally: at which k vertices should we put cones in order to minimize area distortion? The approach we take instead is to penalize the total cone angle—which naturally leads to a convex relaxation.

Consider for instance replacing the problem from Problem 5 with

$$\begin{aligned} & \underset{\mu}{\text{minimize}} && E(u) \\ & \text{subject to} && \Delta u = K \mu \quad \text{in } M \\ & && u = 0 \quad \text{on } \partial M, \end{aligned} \quad (7)$$

where μ is now some kind of smooth curvature distribution, rather than a sum of Dirac deltas. Of course, this problem has a trivial solution $\mu = K$, corresponding to not flattening the surface at all. To obtain nontrivial solutions, we need to somehow encourage μ to be sparse, *i.e.*, to look like a distribution of cones.

L¹ regularization. A tempting idea is to use the L^1 -norm to promote sparsity. The resulting optimization problem is then

$$\begin{aligned} & \underset{f \in L^1(M)}{\text{minimize}} && E(u) + \lambda \|f\|_{L^1} \\ & \text{subject to} && \Delta u = K f \quad \text{in } M \\ & && u = 0 \quad \text{on } \partial M, \end{aligned} \quad (8)$$

where $\lambda > 0$ is a tuning parameter. Though this problem is both convex and seemingly straightforward, it has some serious issues. If one tries to solve it numerically, the result is not a collection of isolated cones: instead, curvature is distributed over larger regions, and the solution jumps around unpredictably for different tessellations of the surface (Fig. 8). Changing, say, the choice of finite elements will not fix this problem, since in the smooth setting, the existence of minimizers to Problem 8 does not hold in general. Hence, even though the discrete, finite dimensional ℓ^1 problem has minimizers, these minimizers do not provide useful solutions. 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 $L^1(M)$. However, the boundedness of the functional does not provide bounds on the derivatives of f and consequently does not imply the L^1 pre-compactness of minimizing sequences.

Measure space regularization. All of these problems can be avoided by replacing the L^1 norm with the measure norm $\|\cdot\|_{\mathcal{M}}$ (Sec. 3.2), and optimizing over measures $\mu \in \mathcal{M}(M)$ (rather than functions):

$$\begin{aligned} & \underset{\mu \in \mathcal{M}(M)}{\text{minimize}} && E(u) + \lambda \|\mu\|_{\mathcal{M}} \\ & \text{subject to} && \Delta u = K \mu \quad \text{in } M \\ & && u = 0 \quad \text{on } \partial M. \end{aligned} \quad (9)$$

For this problem the measure norm concentrates the solution near isolated *points* (*i.e.*, cones) rather than curves or other regions. Although minimizers are in general not exact superpositions of delta masses, they come extremely close—see Sec. 4.5 for further discussion. Most importantly, when we discretize this scheme we obtain sparse cone configurations with very low area distortion (Fig. 8).

4.3 Pre-Dual

When it comes to discretizing Problem 9, we are faced with a challenge: we want to solve a problem involving measures, but standard finite element schemes apply only to functions. Here, a solution is

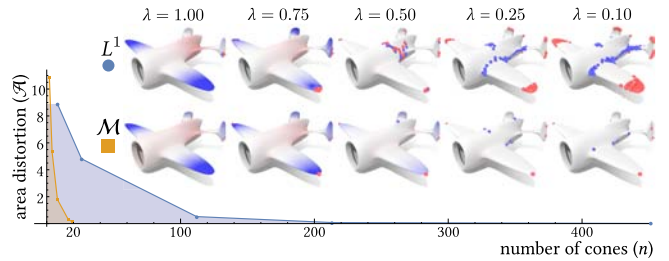


Fig. 8. Using L^1 regularization yields unpredictable results that are not concentrated on isolated points (*top row*), whereas using the measure norm yields well-behaved solutions that properly represent cone singularities (*bottom row*). Tuning the parameter λ does not improve the L^1 result; moreover, the measure-based approach achieves smaller area distortion using dramatically fewer cones (as shown in the plot at bottom).

provided by Fenchel-Rockafellar duality (Sec. 3.4), since the space of measures is the dual of the space of continuous functions (Sec. 3.2). Ultimately, this approach will lead us to a system of optimality equations that we can directly solve for our optimal cone configuration (Sec. 5). A subtle issue is that taking the dual of Problem 9 does *not* yield a problem in terms of functions—however, if we go the other direction and construct a “pre-dual” problem *whose dual* is Problem 9, then we obtain a formulation involving only functions (as desired).

In particular, to express Problem 9 in standard form, we introduce a change of variables $\sigma := K \mu$. Letting $F^*(u)$ be our distortion energy $E(u)$, $G^*(\sigma) = \|K \sigma\|_{\mathcal{M}}$, and $\Lambda^* = \Delta^{-1}$ (or more formally, the solution operator—see App. A.1), the Fenchel dual problem (Problem 3) becomes equivalent to our cone problem, Problem 9. (The reason for introducing the change of variables is that Λ^* must be a linear rather than affine.) One can then show that the corresponding primal problem (Problem 2) is specified by the linear map $\Lambda = \Delta^{-1}$ and the functionals $F : L^2(M) \rightarrow \mathbb{R}$ and $G : C(M) \rightarrow \mathbb{R}$ given by

$$\begin{aligned} F(u) &:= \frac{1}{2} \int_M u^2 dA \\ G(\varphi) &:= \mathbb{1}_\lambda(\varphi) \int_M K\varphi dA, \end{aligned}$$

where $\mathbb{1}_\lambda$ is the indicator function of the λ -ball in $C(M)$, *i.e.*

$$\mathbb{1}_\lambda(\varphi) := \begin{cases} 0 & \text{if } |\varphi(\mathbf{x})| \leq \lambda \text{ for all } \mathbf{x} \in M, \\ +\infty & \text{else.} \end{cases}$$

(See App. A.3.) The resulting primal (or pre-dual) problem then has a very different form from the problem we started with:

$$\begin{aligned} & \underset{\varphi \in C(M)}{\text{minimize}} && \frac{1}{2} \int_M |u|^2 dA \quad \int_M K\varphi dA \\ & \text{subject to} && \Delta \varphi = u \quad \text{in } M, \\ & && \varphi = 0 \quad \text{on } \partial M, \\ & && |\varphi(\mathbf{x})| \leq \lambda \quad \text{for all } \mathbf{x} \in M. \end{aligned} \quad (10)$$

We will refer to the function φ as the *adjoint state*, analogous to a Lagrange multiplier. Notice in particular that this problem no longer involves measures, only ordinary functions u and φ ; the sparsity-inducing measure norm is replaced by inequality constraints, analogous to box constraints that arise in ℓ^1 optimization.

4.4 Optimality System

From here, one idea is to discretize and solve the pre-dual problem directly, but in order to recover the cone configuration we still need the final measure $\mu = K \Delta^2 \varphi$. Numerically, this expression is hard to evaluate directly: for instance, applying the square of the discrete Laplace operator (Sec. 5.1.3) significantly amplifies noise, making it impossible to reliably identify cones. Instead, we first formulate the optimality conditions in the smooth setting, then discretize and solve these optimality conditions directly (Sec. 5). These discrete optimality conditions preserve key relationships between the optimal measure μ and the adjoint variables φ , allowing us to more reliably extract cones. In particular, working out the optimality conditions (App. A.3) and making the change of variables $\bar{\mu} = K \bar{\sigma}$ yields the final optimality system

$$\begin{cases} \Delta \bar{u} = K \bar{\mu}, \\ \Delta \bar{\varphi} = \bar{u}, \\ \bar{\mu} \in \partial \mathbb{1}_\lambda(\bar{\varphi}), \end{cases} \quad (11)$$

where \bar{u} , $\bar{\mu}$, and $\bar{\varphi}$ denote the optimal scale factors, measure, and adjoint state, *resp.*, with zero Dirichlet boundary conditions for \bar{u} and $\bar{\varphi}$. This system will be discretized in Sec. 5.2.1, providing the starting point for our final algorithm.

4.5 Guarantees

What can we say about solutions to our relaxed optimization problem? In what sense are they optimal? Do they solve the original problem—*i.e.*, do we obtain measures that actually represent cones? The basic answer is that we are essentially guaranteed to find cones that minimize L^2 area distortion among all cone configurations of equal (or smaller) total angle. To make this statement more precise, let $\bar{\mu}$ be the solution to the relaxed problem, and consider any other measure $\nu \in \mathcal{M}(M)$ with $\|\nu\|_{\mathcal{M}} \leq \|\bar{\mu}\|_{\mathcal{M}}$; let \bar{u} and ν be the corresponding scale factors, *i.e.*, $\Delta \bar{u} = K \bar{\mu}$ and $\Delta \nu = K \nu$. Then

$$E(\bar{u}) + \lambda \|\bar{\mu}\|_{\mathcal{M}} \leq E(\nu) + \lambda \|\nu\|_{\mathcal{M}},$$

and hence

$$E(\bar{u}) \leq E(\nu) + \lambda(\|\nu\|_{\mathcal{M}} - \|\bar{\mu}\|_{\mathcal{M}}) \leq E(\nu).$$

In other words, the solution $\bar{\mu}$ to the relaxed problem yields minimal area distortion over all curvature measures with norm no greater than $\|\bar{\mu}\|_{\mathcal{M}}$. However, we have not yet established that the optimal measure looks like a sum of Dirac deltas, *i.e.*, that it represents a configuration of cones. The following proposition provides a first step toward understanding the sparsity structure of $\bar{\mu}$:

PROPOSITION 4.1. $\text{supp } \bar{\mu} \subseteq \{\mathbf{x} \in M : |\bar{\varphi}(\mathbf{x})| = \lambda\}$.

If $\bar{\varphi}$ is sufficiently smooth, we can therefore expect that the collection of points where $|\bar{\varphi}| = \lambda$ forms a lower-dimensional set. However, we are able to show that this set *cannot* contain isolated points (Theorem C.4); intuitively, a small curve approximating such a point (such as a small circle around a cone) will always yield a lower-energy solution. In practice we nonetheless find that minimizers are very close to a collection of delta measures; moreover, we can use the theory of optimal transportation to show that rounding such curves onto cones (*à la* Sec. 6.1) will yield an insignificant change in area distortion. In the discrete setting we rarely need to perform

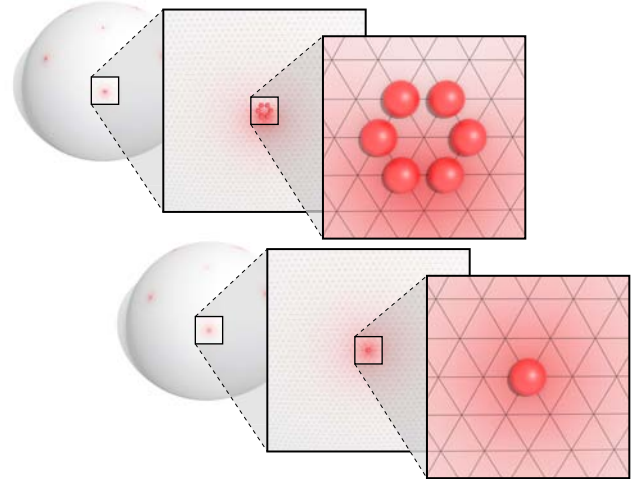


Fig. 9. *Top*: for extremely fine meshes of very smooth surfaces (here, a hemisphere with 256k faces) our algorithm can produce cones arranged in tiny clusters rather than at isolated vertices. *Bottom*: A rigorous stability analysis shows that merging these cones to a nearby vertex (as shown here) cannot yield more than a miniscule change in area distortion, as indicated by red coloring. In practice, such merging is almost never needed.

such rounding since for values of λ in a consistent range (Sec. 6.2), minimizers tend to be supported on isolated vertices—unless the mesh is extremely fine (Fig. 9).

4.6 Moreau-Yosida Regularization

When solving System 11, the rate of convergence will be slow due to the extremely low regularity of minimizers. We therefore consider a sequence of regularized problems which converge to Problem 9 in an appropriate sense, but are much easier to solve. Eventually we can drop regularization entirely and simply solve the original (unregularized) problem, using the most recent regularized solution as an initializer; this technique is known as *Moreau-Yosida regularization*. Note that since the original problem is convex this procedure does not change the minimizer, but rather just serves to improve performance.

For any regularization parameter $\gamma > 0$ consider the problem

$$\begin{aligned} & \underset{\mu_\gamma \in L^2(M)}{\text{minimize}} && E(u_\gamma) + \lambda R(\mu_\gamma) + \frac{\gamma}{2} \|\mu_\gamma\|_{L^2(M)}^2 \\ & \text{subject to} && \Delta u_\gamma = \Omega \mu_\gamma, \end{aligned} \quad (12)$$

where $R(\mu_\gamma) := \|\mu_\gamma dA\|_{\mathcal{M}}$. Since we now consider measures represented by functions, the minimizing scale factors \bar{u}_γ are significantly smoother than \bar{u} . Furthermore, the minimizers $\bar{\mu}_\gamma$ will converge to $\bar{\mu}$ as $\gamma \rightarrow 0$ (in the weak-* sense). Intuitively, as we decrease γ we obtain sharper and sharper approximations of the optimal solution.

Optimality conditions for the regularized problem are given by

$$\begin{cases} \Delta \bar{u}_\gamma = K \bar{\mu}_\gamma, \\ \Delta \bar{\varphi}_\gamma = \bar{u}_\gamma, \\ \bar{\varphi}_\gamma \gamma \bar{\mu}_\gamma \in \partial R(\bar{\mu}_\gamma), \end{cases} \quad (13)$$

again subject to zero Dirichlet boundary conditions.

5 DISCRETIZATION

To discretize our problem, we apply several recent techniques developed in the context of optimal control. In particular, we apply the technique of *variational discretization* to properly treat the discretization of measures [Hinze 2005]; we then use a *semismooth Newton method* to solve the optimality system [Clason and Schiela 2017]—an added bonus here is that, since our system is a principled discretization of a smooth formulation, the rate of convergence is mesh independent [Ulbrich 2011, Chapter 7]. Detailed derivations of our algorithm can be found in App. B; here we focus primarily on the steps needed to actually implement our algorithm. A concise summary of the overall algorithm is given in Fig. 12.

5.1 Preliminaries

5.1.1 Setup. Our algorithm takes as input any manifold triangle mesh $K = (V, E, F)$ with boundary $B \subset V$ (possibly empty). The geometry of this mesh can be specified by vertex positions in \mathbb{R}^3 , though in principle our algorithm really only needs positive edge lengths that determine a piecewise Euclidean metric. The output is a set of *cone vertices* $c_1, \dots, c_k \in V$ and associated *cone angles* ϕ_1, \dots, ϕ_k . (For domains with boundary we may also have boundary data as input or output; see Sec. 8.3.) This data can then be used to compute a parameterization using existing algorithms; for examples in this paper we use a freely-available implementation of *boundary first flattening (BFF)* [Sawhney and Crane 2017]. The user must also provide a parameter $\lambda > 0$ which influences the number of cones (see Sec. 6.2 for further discussion).

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

$$\Omega_i := 2\pi - \sum_{ijk} \theta_i^{jk}. \quad (14)$$

These values are encoded in a column vector $\Omega \in \mathbb{R}^{|I|}$. Similarly, the integrated geodesic curvature (*i.e.*, the geodesic curvature 1-form) associated with a boundary vertex $i \in B$ is discretized via

$$k_i := \pi - \sum_{ijk} \theta_i^{jk}. \quad (15)$$

Again, these values are encoded in a column vector $k \in \mathbb{R}^{|B|}$.

5.1.3 Discrete Poisson Equation. We discretize the Laplace-Beltrami operator Δ via the positive semidefinite *cotan operator* [MacNeal 1949], 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

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

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 ip containing vertex i . We also build a diagonal *lumped mass matrix* $M \in \mathbb{R}^{|V| \times |V|}$:

$$M_{ii} = \frac{1}{3} \sum_{ijk} A_{ijk}. \quad (16)$$

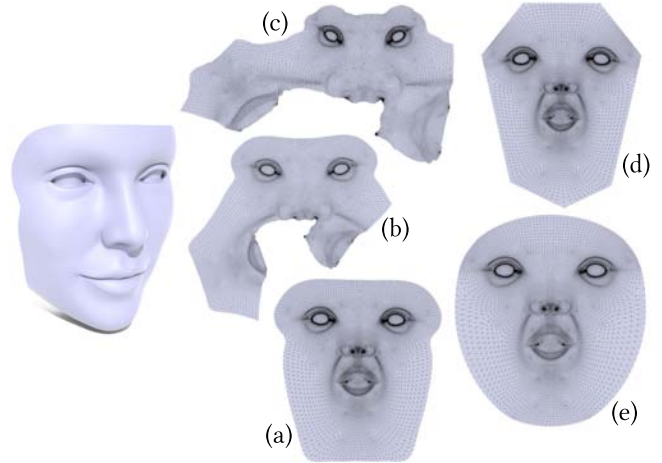


Fig. 10. Different treatments of boundary data allow us to minimize L^2 area distortion with (a) no cones, (b) optimal cones and isometric boundary conditions, (c) optimal cones and optimal boundary conditions; we can also find the least-distorting maps with (d) polygonal and (e) convex boundaries.

Here A_{ijk} is the area of triangle ijk , and we sum over triangles ijk containing vertex i . The matrix equation $Lu = Mf$ then discretizes the Poisson equation $\Delta u = f$ with zero Neumann boundary conditions. More generally, we can partition this system into blocks corresponding to interior vertices (I) and boundary vertices (B) and write

$$\begin{bmatrix} L_{II} & L_{IB} \\ L_{IB}^T & L_{BB} \end{bmatrix} \begin{bmatrix} u_I \\ u_B \end{bmatrix} + \begin{bmatrix} 0 \\ h \end{bmatrix} = \begin{bmatrix} M_{II} & 0 \\ 0 & M_{BB} \end{bmatrix} \begin{bmatrix} f_I \\ f_B \end{bmatrix},$$

which represents a Poisson equation with Neumann boundary conditions $\frac{\partial u}{\partial n} = h$; for each boundary vertex $i \in B$, the value h_i represents half of h integrated over the two boundary edges incident on i . Alternatively, a Poisson equation with Dirichlet boundary conditions $u|_{\partial M} = b$ can be expressed as $L_{II}u_I = M_{II}f_I - L_{IB}b$.

5.2 Derivation of Algorithm

The basic goal of the algorithm is to solve the first-order optimality system (System 11), which we do in two stages:

- **STAGE I** – Solve the 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.

STAGE I produces 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. As mentioned in Sec. 4.6, this approach substantially accelerates convergence relative to solving the unregularized system directly—in practice we observe about an order of magnitude speedup.

5.2.1 Discrete Optimality System. To discretize System 11, one first replaces the Laplace-Beltrami operator by the cotan-Laplacian

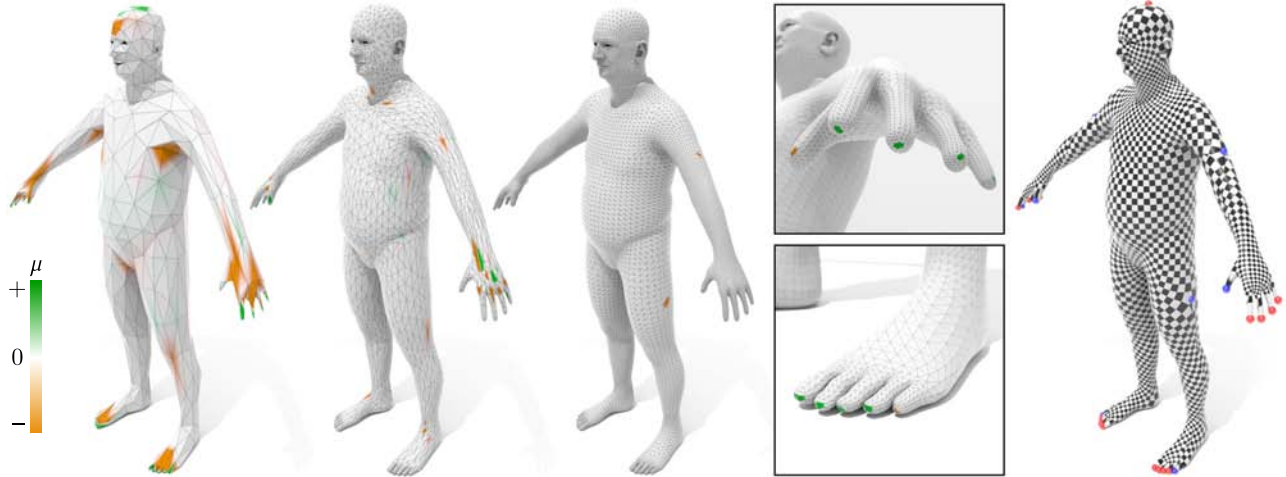


Fig. 11. Mesh hierarchy built on the *Old Man Multires* model; here we use only three levels. Compared to solving directly on the full-resolution 30k triangle mesh, we obtain a speedup factor of about 23x (from 14.77s to about 0.62s). Colors indicate the optimal regularized measure; in the final (fine) 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 .

to obtain the system

$$\begin{cases} L_{\parallel} \mathbf{u} = \Omega - \mu, \\ L_{\parallel}^{\top} \varphi = M_{\parallel} \mathbf{u}, \\ \mu \in \partial \mathbb{1}_{\lambda}(\varphi), \end{cases} \quad (17)$$

Discretizing the third optimality condition is not as straightforward, but can ultimately be achieved via a simple projection operator. In particular, let

$$P_{\lambda}(\mathbf{v}) := \max(0, \mathbf{v} - \lambda) + \min(0, \mathbf{v} + \lambda), \quad (18)$$

where max and min are applied componentwise. As shown in App. B.2, the relationship $\mu \in \partial \mathbb{1}_{\lambda}(\varphi)$ is then equivalent to the condition

$$\mu = P_{\lambda}(\varphi + \mu). \quad (19)$$

We can use this equivalence to express discrete optimality in terms of a function $F : \mathbb{R}^{3|I|} \rightarrow \mathbb{R}^{3|I|}$ given by

$$F(\mathbf{u}, \varphi, \mu) := \begin{bmatrix} L_{\parallel} \mathbf{u} - \Omega + \mu \\ L_{\parallel}^{\top} \varphi - M_{\parallel} \mathbf{u} \\ \mu - P_{\lambda}(\varphi + \mu) \end{bmatrix}. \quad (20)$$

We then seek solutions to the system of equations $F(\mathbf{u}, \varphi, \mu) = 0$. Since this map is both nonlinear and nonsmooth (due to the projection P_{λ}), we solve it using a semismooth Newton method (as described in Sec. 5.3).

5.2.2 Regularized Problems. As discussed in Sec. 4.6, we solve a sequence of regularized problems whose solutions ultimately converge to the minimizer. In particular, we discretize System 13 as

$$\begin{cases} L_{\parallel} \mathbf{u} = \Omega - \mu, \\ L_{\parallel}^{\top} \varphi = M_{\parallel} \mathbf{u}, \\ \mu = \frac{1}{\gamma} P_{\lambda}(\varphi). \end{cases} \quad (21)$$

(see App. B.2 for a derivation). We can then substitute the expression for μ into the first equation to obtain a system involving only the

variables $\mathbf{u}, \varphi \in \mathbb{R}^{|I|}$. We therefore seek solutions to the system of equations $F_{\gamma}(\mathbf{u}, \varphi) = 0$, where $F_{\gamma} : \mathbb{R}^{2|I|} \rightarrow \mathbb{R}^{2|I|}$ is given by

$$F_{\gamma}(\mathbf{u}, \varphi) := \begin{bmatrix} L_{\parallel} \mathbf{u} - \Omega + \frac{1}{\gamma} P_{\lambda}(\varphi) \\ L_{\parallel}^{\top} \varphi - M_{\parallel} \mathbf{u} \end{bmatrix}. \quad (22)$$

5.3 Semismooth Newton Method

We apply a semismooth Newton method to solve the optimality systems encoded by Eqn. 20 and 22, which generalizes Newton's method to functions that are not classically (*i.e.*, Fréchet) differentiable. Like the standard Newton method, the semismooth version aims to find zeros of a function $F(\mathbf{x}) = 0$ by iteratively computing zeros of a linearized version. Each step, the current iterate \mathbf{x} is updated by solving the system

$$DF(\mathbf{x})\mathbf{y} = -F(\mathbf{x}), \quad (23)$$

and applying the update $\mathbf{x} \leftarrow \mathbf{x} + \mathbf{y}$. The main difference is that the linear map $DF(\mathbf{x})$ is the *semismooth differential* at \mathbf{x} , rather than the Jacobian (which may not be well-defined). The semismooth differential effectively determines the descent direction, and must be tailored to the given function F (see Sec. 5.3.2). The process then continues until some specified termination conditions are satisfied (Sec. 5.3.3). In principle one can guarantee global convergence of such methods by adopting a trust region strategy [Ulbrich 2011, Chapter 5], though in practice we find that the basic semismooth Newton strategy is always sufficient, perhaps due to our use of a regularizing sequence.

5.3.1 Active Sets. To build the semismooth differential for our problem (and to check for convergence), we will need to keep track of the *active sets* where either the max or min operators in the projection P_{λ} are nonzero. These sets can be specified by binary values at vertices—in particular, for any given argument $\mathbf{x} \in \mathbb{R}^{|I|}$ let

$A^+(x)$ and $A^-(x)$ be the $|| \times ||$ diagonal matrices with entries

$$\begin{aligned} A_{ii}^+(x) &:= x_i > +\lambda, \\ A_{ii}^-(x) &:= x_i < -\lambda \end{aligned} \quad (24)$$

for each vertex $i \in I$. In **STAGE I**, the active sets are then given by $A^+(\varphi)$ and $A^-(\varphi)$, and in **STAGE II** they are given by $A^+(\varphi + \mu)$ and $A^-(\varphi + \mu)$.

5.3.2 Semismooth Differentials. To apply this approach, we need a semismooth differential for F and F_γ . At a point $x \in \mathbb{R}^{|V|}$ a semismooth differential of P_λ is given by the matrix $D(x) := A^+(x) + A^-(x)$ (see [Hintermüller et al. 2002, Lemma 3.1]). Since the remaining operators in F and F_γ are linear, the overall semismooth differentials are given by

$$DF_\gamma(u, \varphi) := \begin{bmatrix} L_{||} & \frac{1}{\gamma} D(\varphi) \\ -M_{||} & L_{||}^T \end{bmatrix} \quad (25)$$

and

$$DF(u, \varphi, \mu) := \begin{bmatrix} L_{||} & 0 & I \\ -M_{||} & L_{||}^T & 0 \\ 0 & -D(\varphi + \mu) & I - D(\varphi + \mu) \end{bmatrix}, \quad (26)$$

where 0 and I denote the $|| \times ||$ zero and identity matrix, *resp.*

5.3.3 Convergence Criteria. To check convergence, we simply check if either the norm of the current residual b is below a small tolerance $\varepsilon > 0$, or if the active sets are no longer changing, *i.e.*, in **STAGE I** we check whether $A^+(\varphi)$ and $A^-(\varphi)$ were changed by the most recent Newton step; in **STAGE II** we instead check $A^+(\varphi + \mu)$ and $A^-(\varphi + \mu)$. An analysis of this stopping criteria can be found in Hintermüller et al. [2002].

5.3.4 Boundary Control. Here, we discuss the practical changes to the algorithm to incorporate Dirichlet boundary conditions $b \in \mathbb{R}^{|B|}$ —a comprehensive discussion and derivation of the algorithm with Dirichlet boundary conditions is presented in App. A.1.1 and App. A.2. To incorporate Dirichlet boundary values into the optimization we replace the operator $F(u, \varphi, \mu)$ with

$$F(u, \varphi, d, \mu, b) := \begin{bmatrix} L_{||}u - \Omega - \mu - L_{IB}b \\ \mu - P_\lambda(\varphi + \mu) \\ L_{||}^T\varphi - M_{||}u - L_{IB}d \\ d - P_\lambda(d + b), \end{bmatrix}.$$

The vectors $b, d \in \mathbb{R}^{|B|}$ encode the Dirichlet boundary data for the original and pre-dual problems, *resp.* Similarly, to optimize over the Neumann boundary values $h \in \mathbb{R}^{|B|}$ (which are the same for the original and pre-dual problems) we consider the optimality system

$$F(u, \varphi, \mu, h) := \begin{bmatrix} L_{||}u_I + L_{IB}u_B - \Omega - \mu \\ L_{IB}^T u_I + L_{BB}u_B - k - h \\ L_{||}\varphi_I + L_{IB}\varphi_B - M_{||}u_I \\ L_{BI}^T\varphi_B + L_{BB}\varphi_B - M_{BB}u_B - h \\ \mu - P_\lambda(\varphi + \mu) \end{bmatrix}.$$

To derive the semismooth differential one then proceeds exactly as in the case without boundary, *i.e.*, use the semismooth differential for P_λ ; use the ordinary differential for all other terms.

STAGE I

- $u \leftarrow 0 \in \mathbb{R}^{|I|}, \varphi \leftarrow 0 \in \mathbb{R}^{|E|}, \gamma \leftarrow 1$
- **for** $\gamma = 10^0, 10^{-1}, 10^{-2}, \dots, 10^{-N}$:
 - $(u, \varphi) \leftarrow \text{COMPUTEREGULARIZEDMEASURE}(u, \varphi, \lambda, \gamma)$

STAGE II

- $(u, \varphi, \mu) \leftarrow \text{COMPUTEOPTIMALMEASURE}(u, \varphi, \lambda)$

Extract final cones from μ (Sec. 6.1)

COMPUTEREGULARIZEDMEASURE

- **Until** convergence (Sec. 5.3.3):
 - Evaluate active sets (Eqn. 24)
 - Evaluate residual $b_I := F_\gamma(u, \varphi)$ (Eqn. 22)
 - Evaluate differential $U_I := DF_\gamma(u, \varphi)$ (Eqn. 25)
 - Solve $U_I[v^T q^T]^T = -b_I$
 - $u \leftarrow u + v$
 - $\varphi \leftarrow \varphi + q$

COMPUTEOPTIMALMEASURE

- $\mu \leftarrow P_\lambda(\varphi)$ (Eqn. 18)
- **Until** convergence (Sec. 5.3.3):
 - Evaluate active sets (Eqn. 24)
 - Evaluate residual $b_{II} := F(u, \varphi, \mu)$ (Eqn. 20)
 - Evaluate differential $U_{II} := DF(u, \varphi, \mu)$ (Eqn. 26)
 - Solve $U_{II}[v^T q^T v^T]^T = -b_{II}$
 - $u \leftarrow u + v$
 - $\varphi \leftarrow \varphi + q$
 - $\mu \leftarrow \mu + v$

Fig. 12. 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.

6 ALGORITHMIC CONSIDERATIONS

6.1 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 equal to a cone angle $\phi_{c_i} \gg 0$, or they are numerically zero—we use a tolerance of 10^{-12} . Very rarely cones may appear in tiny clusters, reflecting the fact that in the smooth setting one can slightly reduce area distortion by replacing a Dirac measure at a point p with a measure supported on a tiny ring around p (Sec. 4.5). 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. A stability result shows that this rounding procedure cannot change the area distortion by more than a tiny amount (Theorem C.5).

6.2 Tuning Parameter

The parameter $\lambda \geq 0$ influences the number of cone singularities, or more precisely, the maximum allowable total cone angle $\Phi = \sum_i |\phi_i|$. Decreasing λ reduces the distortion at the cost of greater total cone angle, and vice versa. As with many recent methods (*e.g.*, [Myles

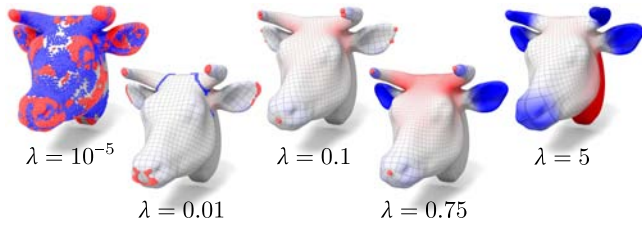


Fig. 13. The parameter λ controls the strength of the penalty on the total cone angle, and in turn, influences the number of cones. For values above 1 (strong penalty) we tend to see that no cones are placed. For values very close to zero (weak penalty) the curvature measure stops being sparse, and we get cones with small angles placed densely across regions or along curves rather than at isolated points.

and Zorin 2012, 2013; Knöppel et al. 2013]) the relationship between algorithm parameters and the behavior of singularities is indirect. However, we find that the effect of λ is surprisingly consistent across a wide variety of different meshes and tessellations (see Figures 14 and 18), as long as we normalize it by the total surface area—we make this assumption throughout. In practice, we therefore do not need to do extensive parameter tuning: for instance, universally setting $\lambda = 1/2$ already provides an automatic way to get reasonable cone configurations (Fig. 15). 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. We find that a good range of values across a wide variety of examples at different resolutions is $\frac{1}{10} < \lambda < 1$ (see Figures 13 and 14). 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. Some values in the range $0 < \lambda < \frac{1}{10}$ yield configurations of cones arranged along curves. (See Fig. 13.) Finally, since solutions are consistent across different levels of tessellation (see Fig. 18, top), one could quickly tune this parameter on a coarse mesh before computing the fine solution.

6.3 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 K we first construct a sequence of progressively finer meshes $K_1, \dots, K_M = K$, 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 K_l , we identify the set of vertices on the next finest mesh K_{l+1} that are closer to i than any other vertex in K_l . The values of u_i and φ_i are then equally distributed over these vertices. Note that one does

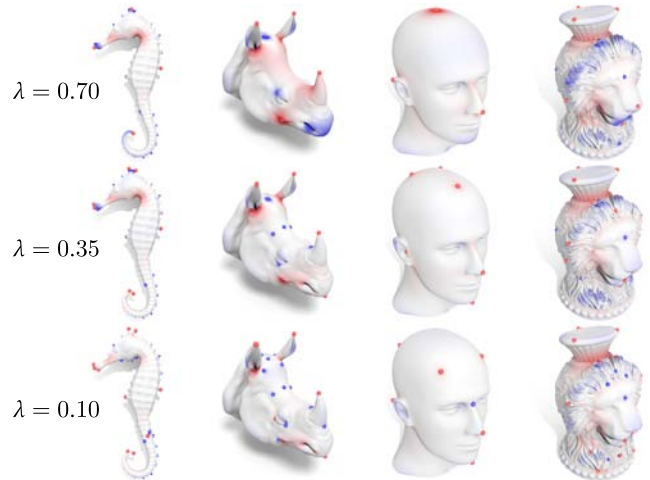


Fig. 14. The parameter λ has a very consistent behavior across different surfaces, typically producing a similar density of cones.

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* [Schmidt and Singh 2010], and use a constant $s = 2$. Fig. 11 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 paper were computed directly on the fine input mesh.

7 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. Since the formulation is convex, we obtain identical results for any initialization. The multiresolution strategy outlined in Sec. 6.3 reduces the size of these systems substantially, though we did not find it essential for most of the examples in this paper: for models of about 100–150k triangles the algorithm takes at most about 20–25 seconds. We did little optimization of our code; there are plenty of opportunities for acceleration in terms of both linear algebra and numerical algorithms.

7.1 Comparisons

We here compare the results of our method (**MAD**) to existing cone singularity placement strategies introduced by Ben-Chen et al. [2008] (**CPMS**) and Myles and Zorin [2012] (**GPIF**). The basic cone placement strategy from Springborn et al. [2008] (**CETM**) is similar to **CPMS**, but we omit a comparison since their strategy for picking angles is highly mesh dependent as discussed in Sec. 4.1.2. Likewise, we do not compare to the recent method of Vintescu et al. [2017a] which provides only cone *angles* and not the number of cones or their positions, nor the method of Myles and Zorin [2013] which

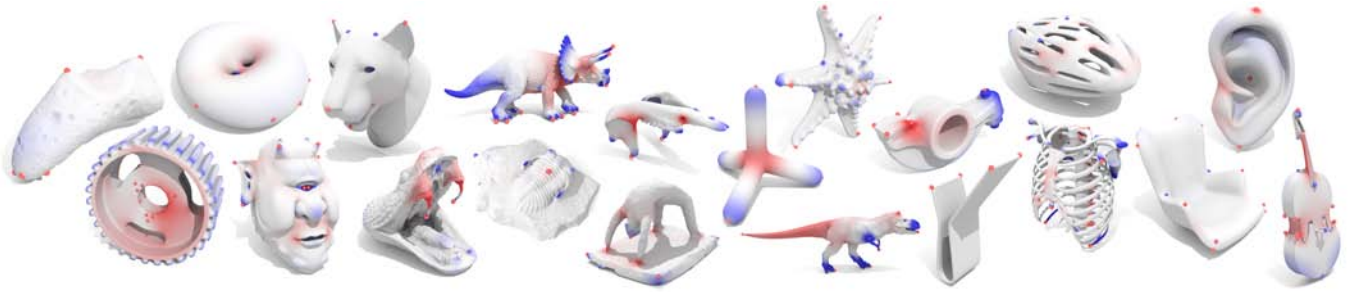


Fig. 15. Using a default parameter $\lambda = 1/2$ allows one to flatten a wide variety of different models without any explicit tuning or adjustment.

provides very similar results to **GPIF** in the absence of quantization and field alignment ([Myles and Zorin 2013, Fig. 14]). **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.

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 (App. C) 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. 24, *top*). In other examples, we obtain much lower area distortion, or alternatively, comparable distortion with far fewer cones (see for instance Figures 3 and 16, and 24, *bottom left*). In Fig. 16, we also see that lines of singularities (as sometimes placed by **GPIF**) do not necessarily yield lower area distortion than simply placing a few carefully-selected cones.

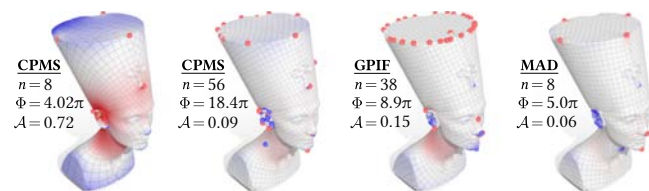


Fig. 16. Finding *just* the right configuration of cones and angles can sometimes dramatically reduce area distortion. Here, **MAD** almost completely eliminates area distortion using just 8 cones (*far right*). Using the same number of cones in **CPMS** (*far left*) yields far greater distortion; alternatively, 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.

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 .

7.2 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. 18, 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. 20). 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. 19 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. 4.1.2 (**GPIF** does not suffer from this same artifact).

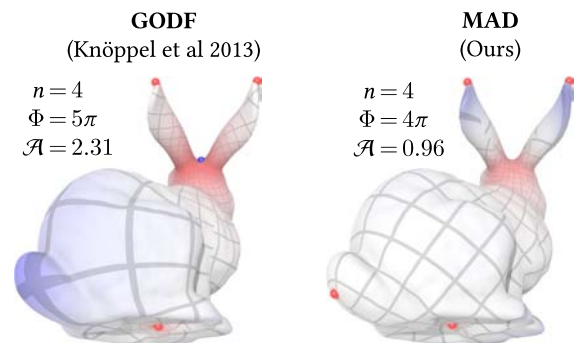


Fig. 17. Optimal singularities for direction fields (*top*) are not necessarily good for conformal flattening (*bottom*).

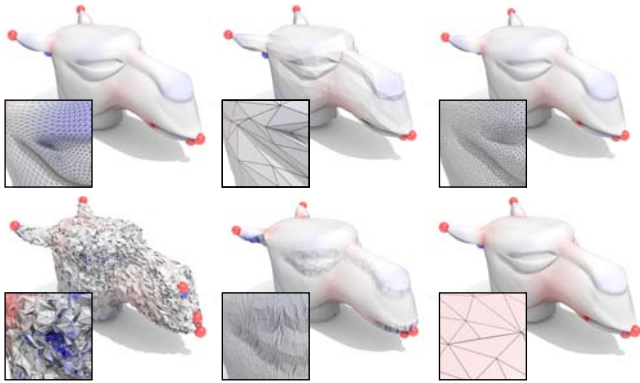


Fig. 18. Our method produces consistent results on meshes of very different resolution (*top row*) and is also robust to meshing artifacts such as noise (*bottom left*), anisotropy (*bottom center*), and severely non-Delaunay elements (*bottom right*). The same λ value is used in all examples.

8 EXTENSIONS

Our basic optimization framework is flexible enough to be extended in a variety of ways—a precise formulation of these extensions is detailed in App. A.2.

8.1 Nonuniform Importance

We can augment our method to influence both (i) where cones are placed and (ii) where distortion is measured. As detailed in App. A.2, one can either provide continuous functions $w_{\mathcal{G}}, w_{\mathcal{R}} : V \rightarrow \mathbb{R}_{>0}$ that act as a penalty on distortion and cone placement (*resp.*), or binary functions $U_E, U_R : V \rightarrow \{0, 1\}$ that explicitly excludes regions where distortion is measured and cones are placed (*resp.*). The functions U_E, U_R are 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. 8.3 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 an intrinsic local feature size at each vertex $i \in V$ (*à la* [Sharp and Crane 2018,

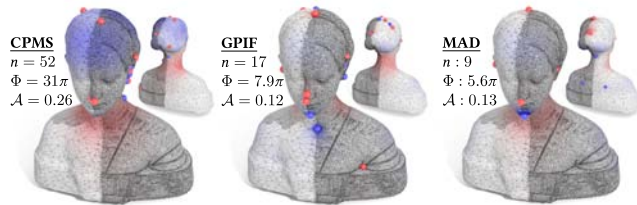


Fig. 19. Variable density mesh. *Left*: CPMS places far more singularities in finely tessellated regions, where Green’s functions are better resolved; note that many spheres overlap due to close clustering of cones. *Center*: GPIF also violates symmetry, and achieves lower distortion than MAD only by 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.

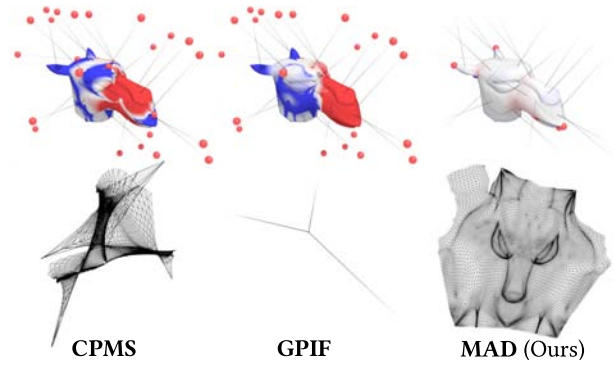


Fig. 20. Stress test of robustness. Methods that place cones according to peaks in scale factors (*left*) or based on curvature (*center*) may be confounded by outliers; here we consider an extreme case where near-invisible spikes are added to a mesh, leading to cone configurations that are impossible to parameterize (*bottom*). Since our method minimizes an integral energy, it is generally not confounded by outliers or noise, in this case ignoring the spikes and leading to a high-quality parameterization (*right*).

Sec. 3.5.1]), *i.e.*, a value $r_i := 1/(|\Omega_i|/M_{\parallel} + \varepsilon)$, where Ω_i is the angle defect (Eqn. 14) and M_{\parallel} is the barycentric dual area. This value will be small in flat regions and large in highly curved regions. We then set $(w_{\mathcal{G}})_i := r_i$ 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. 7.

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. 21); here we likewise need only penalize *distortion* in the visible regions. Given a particular viewpoint, we set $(U_E)_i$ to 1 and $(U_R)_i$ 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 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. We are now free to consider 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.

8.2 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,




Fig. 21. 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*).

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. 22, (*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. 8.1). In another case, we simply optimize over all nonnegative cone configurations with total angle 4π , *without* including any kind of sparsity-inducing norm—amazingly enough, we still get a sparse solution.

Another example where angle bounds are potentially useful is in finding cone configurations for seamless *integer grid maps* [Bommes et al. 2013a], 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. 23 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.

8.3 Optimizing the Boundary

Incorporating boundary data into our optimization problem provides additional control, as illustrated in Fig. 10 (see App. A.1.1 for details). For conformal flattenings without cone singularities, Springborn et al. [2008, App. E] show that minimal area distortion (with respect to E_D) is achieved by constant Dirichlet boundary conditions $u|_{\partial M} \equiv \text{const}$. However, these boundary conditions are not necessarily optimal in the presence of cones, nor when minimizing the energy E . We therefore augment Problem 9 to jointly optimize both the cone configuration and the choice of Dirichlet boundary

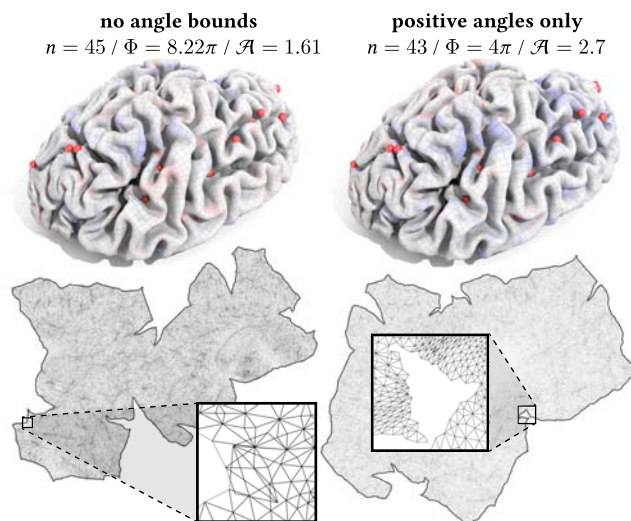


Fig. 22. *Left*: When finally mapped to the plane a cone flattening of a surface (such as this brain) may have local noninjectivity at negative cones, unless these cones are cut into sufficiently small pieces (see zoom). *Right*: finding an optimal solution with only positive cones avoids this source of local noninjectivity.

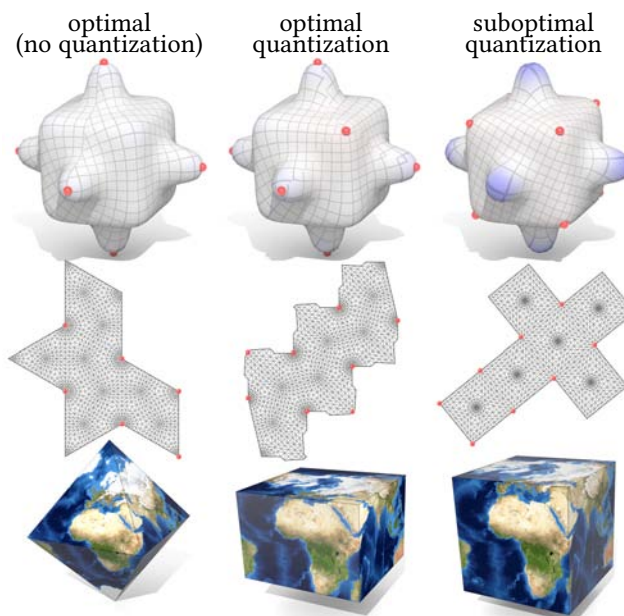



Fig. 23. *Left*: in the absence of any bounds on cone angles ϕ_i , an optimal configuration for this model is to place eight equal cones corresponding to an octahedron. *Center*: if we now limit angles to the range $\phi \in [0, \pi/2]$, we get a configuration that has eight cones of $\pi/2$, but *not* at the corners of a cube as one might expect: instead, we get a flattening to a cuboid with unequal lengths. *Right*: the more intuitive configuration with cones at cube corners yields higher area distortion.

conditions (App. A.1.1), or to just optimize over one or the other. Alternatively, we can get polygonal boundaries by promoting the sparsity of boundary curvature, achieved by penalizing the measure norm of Neumann boundary data (consider the boundary conditions in Eqn. 1). Finally, forcing boundary curvature to be positive (*à la* Sec. 8.2) leads to minimally-distorting flattenings with *convex* boundaries.

9 LIMITATIONS AND FUTURE WORK

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 (App. C) that leads to a practical rounding procedure in the very rare case where cones appear in smaller clusters (Sec. 6.1). 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. 8.1 (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. 14). 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 [Del Pino and Román 2015; De Marchis and López-Soriano 2016]. 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. One nice feature of cone parameterization is that the flattening is performed intrinsically, prior to the final 2D layout—an interesting open question is how to find the *best* 2D layout, *e.g.*, the one that yields greatest signal quality (Fig. 25). 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. The analytical perspectives developed here may provide new ways of looking at these problems.



ACKNOWLEDGEMENTS

Thanks to Henrik Schumacher for useful references to algorithms from optimal control, to Rohan Sawhney for implementation help with BFF, and to Alex Huth for sharing cortical surface data [Gao et al. 2015]. This work was sponsored in part by NSF Awards CCF 1717320 and DMS 1516677, a CMU Summer Undergraduate Research Fellowship (SURF), and gifts from Autodesk Research and Adobe Research.

REFERENCES

- Noam Aigerman, Shahar Z. Kovalsky, and Yaron Lipman. 2017. Spherical Orbifold Tutte Embeddings. *ACM Trans. Graph.* 36, 4 (2017).
- Noam Aigerman and Yaron Lipman. 2015. Orbifold Tutte Embeddings. *ACM Trans. Graph.* 34, 6 (2015).
- Noam Aigerman and Yaron Lipman. 2016. Hyperbolic Orbifold Tutte Embeddings. *ACM Trans. Graph.* 35, 6 (2016).
- Luigi Ambrosio, Edoardo Mainini, and Sylvia Serfaty. 2011. Gradient flow of the Chapman–Rubinstein–Schatzman model for signed vortices. In *Annales de l’Institut Henri Poincaré Non Linear Analysis*, Vol. 28.
- Thierry Aubin. 1998. *Some nonlinear problems in Riemannian geometry*. Springer Science & Business Media.
- Heinz H Bauschke and Patrick L Combettes. 2017. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer.
- Mirela Ben-Chen, Craig Gotsman, and Guy Bunin. 2008. Conformal flattening by curvature prescription and metric scaling. In *Computer Graphics Forum*, Vol. 27. 449–458.
- David Bommes, Marcel Campen, Hans-Christian Ebke, Pierre Alliez, and Leif Kobbelt. 2013a. Integer-grid Maps for Reliable Quad Meshing. *ACM Trans. Graph.* 32, 4 (2013).
- David Bommes, Bruno Lévy, Nico Pietroni, Enrico Puppo, Claudio Silva, Marco Tarini, and Denis Zorin. 2013b. Quad-Mesh Generation and Processing. (2013).
- Haim Brezis. 2010. *Functional analysis, Sobolev spaces and partial differential equations*. Springer Science & Business Media.
- Eduardo Casas, Christian Clason, and Karl Kunisch. 2012. Approximation of elliptic control problems in measure spaces with sparse solutions. *SIAM Journal on Control and Optimization* 50, 4 (2012).
- Isaac Chao, Ulrich Pinkall, Patrick Sanan, and Peter Schröder. 2010. A Simple Geometric Model for Elastic Deformations. *ACM Trans. Graph.* 29, 4 (2010).
- Xiaojun Chen, Zuhair Nashed, and Liqun Qi. 2000. Smoothing methods and semismooth methods for nondifferentiable operator equations. *SIAM J. Numer. Anal.* 38, 4 (2000).
- Christian Clason and Karl Kunisch. 2012. A measure space approach to optimal source placement. *Computational Optimization and Applications* 53, 1 (2012).
- Christian Clason and Anton Schiela. 2017. Optimal control of elliptic equations with positive measures. *ESAIM: Cont., Opt. and Calc. of Var.* 23, 1 (2017).
- Teresa D’Aprile, Francesca De Marchis, and Isabella Ianni. 2016. Prescribed Gauss curvature problem on singular surfaces. *arXiv preprint arXiv:1612.03657* (2016).
- Francesca De Marchis and Rafael López-Soriano. 2016. Existence and non existence results for the singular Nirenberg problem. *Calculus of Variations and Partial Differential Equations* 55, 2 (2016).
- Manuel Del Pino and Carlos Román. 2015. Large conformal metrics with prescribed sign-changing Gauss curvature. *Calculus of Variations and Partial Differential Equations* 54, 1 (2015).
- Jeff Erickson and Sariel Har-Peled. 2004. Optimally cutting a surface into a disk. *Discrete & Computational Geometry* 31, 1 (2004).
- James S. Gao, Alexander G. Huth, Mark D. Lescroart, and Jack L. Gallant. 2015. Pycortex: an interactive surface visualizer for fMRI. *Front. in Neuro.* 9 (2015).
- Andreas Günther and Moulay Hicham Tber. 2016. A goal-oriented adaptive Moreau-Yosida algorithm for control-and state-constrained elliptic control problems. *Advances in Applied Mathematics and Mechanics* 8, 3 (2016).
- Lester La Verne Helms. 2009. *Potential theory*. Springer.
- Heinrich Hencky. 1928. Über die Form des Elastizitätsgesetzes bei ideal elastischen Stoffen. *Zeitschrift für technische Physik* 9 (1928).
- M. Hintermüller, K. Ito, and K. Kunisch. 2002. The Primal-Dual Active Set Strategy As a Semismooth Newton Method. *SIAM J. on Optimization* 13, 3 (2002).
- Michael Hinze. 2005. A Variational Discretization Concept in Control Constrained Optimization: The Linear-Quadratic Case. *Comp. Opt. and Appl.* 30, 1 (2005).
- Kazufumi Ito and Karl Kunisch. 2003a. Semi-smooth Newton methods for state-constrained optimal control problems. *Systems & Control Letters* 50, 3 (2003).
- Kazufumi Ito and Karl Kunisch. 2003b. Semi-smooth Newton methods for variational inequalities of the first kind. *ESAIM: Math. Model. and Num. Anal.* 37, 1 (2003).
- Liliya Kharevych, Boris Springborn, and Peter Schröder. 2006. Discrete Conformal Mappings via Circle Patterns. *ACM Trans. Graph.* 25, 2 (2006).

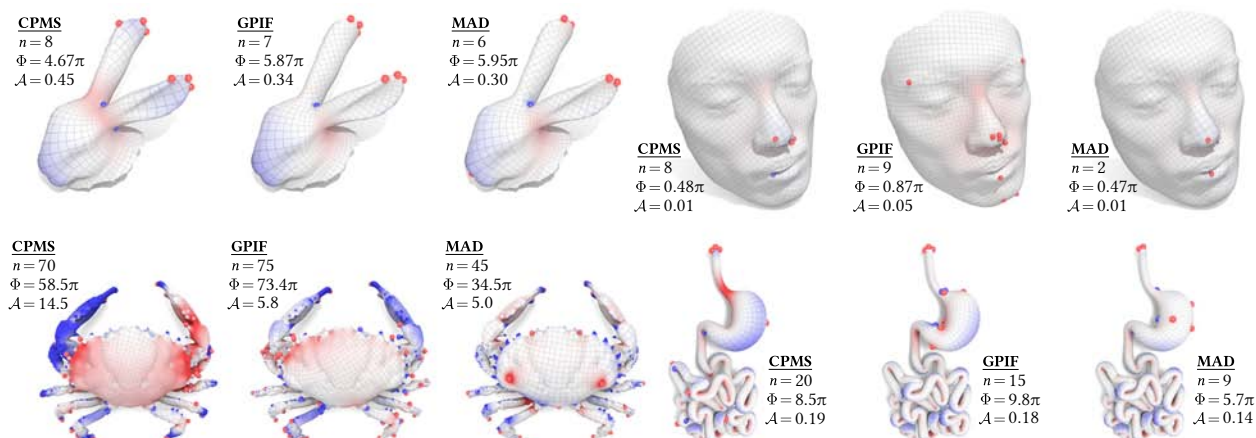


Fig. 24. More comparisons. On models with simple geometry (*top row*) greedy or region-growing strategies can work quite well, though **MAD** still performs slightly better. On more challenging models such as the crab (*bottom left*) the gap typically widens—note also the high degree of symmetry for **MAD**.

Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2013. Globally optimal direction fields. *ACM Trans. Graph.* 32, 4 (2013).

Mina Konakovic, Keenan Crane, Bailin Deng, Sofien Bouaziz, Daniel Piker, and Mark Pauly. 2016. Beyond Developable: Computational Design and Fabrication with Auxetic Materials. *ACM Trans. Graph.* 35, 4 (2016).

Walter Littman, Guido Stampacchia, and Hans F Weinberger. 1963. Regular points for elliptic equations with discontinuous coefficients. *Annali della Scuola Normale Superiore di Pisa-Classe di Scienze* 17, 1-2 (1963).

Richard MacNeal. 1949. *The Solution of Partial Differential Equations by Means of Electrical Networks*. Ph.D. Dissertation. California Institute of Technology.

E. Mainini. 2012. A description of transport cost for signed measures. *Journal of Mathematical Sciences* 181 (2012).

Haggai Maron, Meirav Galun, Noam Aigerman, Miri Trope, Nadav Dym, Ersin Yumer, Vladimir G. Kim, and Yaron Lipman. 2017. Convolutional Neural Networks on Surfaces via Seamless Toric Covers. *ACM Trans. Graph.* 36, 4 (2017).

Ashish Myles and Denis Zorin. 2012. Global parametrization by incremental flattening. *ACM Transactions on Graphics (TOG)* 31, 4 (2012).

Ashish Myles and Denis Zorin. 2013. Controlled-distortion constrained global parametrization. *ACM Transactions on Graphics (TOG)* 32, 4 (2013).

Herzog Roland and Kunisch Karl. 2010. Algorithms for PDE-constrained optimization. *GAMM-Mitteilungen* 33, 2 (2010), 163–176.

E.B. Saff and V. Totik. 1997. *Logarithmic Potentials with External Fields*. Springer Berlin Heidelberg.

Rohan Sawhney and Keenan Crane. 2017. Boundary First Flattening. *ACM Trans. Graph.* 37, 1 (2017).

Ryan Schmidt and Karan Singh. 2010. Meshmixer: An Interface for Rapid Mesh Composition. In *ACM SIGGRAPH 2010 Talks (SIGGRAPH '10)*.

Nick Sharp and Keenan Crane. 2018. Variational Surface Cutting. *ACM Trans. Graph.* 37, 4 (2018).

Alla Sheffer, Emil Praun, and Kenneth Rose. 2006. Mesh Parameterization Methods and Their Applications. *Found. Trends. Comput. Graph. Vis.* 2, 2 (2006).

Yousuf Soliman. 2018. *Conformal Cone Parameterization through Optimal Control*. Master's thesis. Carnegie Mellon University.

Olga Sorkine and Marc Alexa. 2007. As-rigid-as-possible Surface Modeling. In *Proc. Symp. Geom. Proc. (SGP '07)*.

Boris Springborn, Peter Schröder, and Ulrich Pinkall. 2008. Conformal equivalence of triangle meshes. In *ACM Transactions on Graphics (TOG)*, Vol. 27.

Guido Stampacchia. 1965. Le problème de Dirichlet pour les équations elliptiques du second ordre à coefficients discontinus. *Ann. Inst. Fourier* 15, 1 (1965).

William Thurston. 2002. *The Geometry and Topology of 3-Manifolds*. (2002).

Marc Troyanov. 1989. Metrics of constant curvature on a sphere with two conical singularities. *Lect. Notes Math* 1410 (1989).

Marc Troyanov. 1991. Prescribing curvature on compact surfaces with conical singularities. *Trans. Amer. Math. Soc.* 324, 2 (1991).

Alex Tsui, Devin Fenton, Phong Vuong, Joel Hass, Patrice Koehl, Nina Amenta, David Coeurjolly, Charles DeCarli, and Owen Carmichael. 2013. Globally Optimal Cortical Surface Matching With Exact Landmark Correspondence. In *Information Processing in Medical Imaging*.

M. Ulbrich. 2002. *Nonsmooth Newton-like Methods for Variational Inequalities and Constrained Optimization Problems in Function Spaces*. Habilitationsschrift. Fakultät für Mathematik, Technische Universität München, Germany.

Michael Ulbrich. 2011. *Semismooth Newton Methods for Variational Inequalities and Constrained Optimization Problems in Function Spaces*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.

Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional Field Synthesis, Design, and Processing. *Computer Graphics Forum* (2016).

Ana Vintescu, Florent Dupont, Guillaume Lavoué, Pooran Memari, and Julien Tierny. 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*.

Ana-Maria Vintescu, Florent Dupont, Guillaume Lavoué, Pooran Memari, and Julien Tierny. 2017b. Conformal Factor Persistence for Fast Hierarchical Cone Extraction. In *Eurographics Short Papers*.

Eugene Zhang, Konstantin Mischaikow, and Greg Turk. 2005. Feature-based Surface Parameterization and Texture Mapping. *ACM Trans. Graph.* 24, 1 (2005).

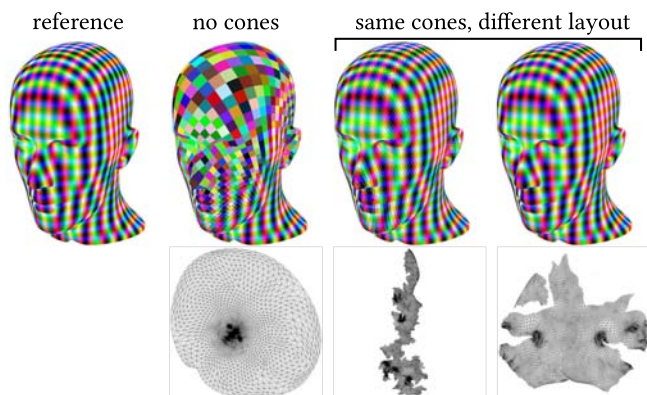


Fig. 25. Reducing area distortion improves signal quality. Following Zhang et al. [2005], a given signal (*far left*) is sampled onto 2D image, which is then displayed on the surface. Without cone singularities (*center left*) the signal is badly undersampled in regions of large scale distortion. Incorporating cones dramatically improves sampling (*center and far right*), though finding the *best* 2D layout for a given cone metric remains an interesting open question.

A SMOOTH FORMULATION (DETAILS)

In this appendix, we present the precise mathematical framework which we use to employ Fenchel-Rockafellar duality. In particular, we discuss the solution operator to the Poisson equation with measure data and present all of the computations and details necessary to apply convex duality.

A.1 Poisson Equation with Measure Data

The PDE constraint present in our relaxed formulation Problem 9 is a simple Poisson equation

$$\Delta u = \mu,$$

where $\mu \in \mathcal{M}(M)$ is a measure. Since M is a two-dimensional manifold the classical variational framework is inequipt to deal with measure valued right hand sides. Instead, we rely on the *harmonic Green's function* to obtain solutions:

THEOREM A.1 (AUBIN [1998, THEOREM 4.17]). *Let \bar{M} be a smooth manifold with boundary ∂M of class C^∞ . Then, there exists $G_M : \bar{M} \times \bar{M} \rightarrow \mathbb{R}$, which is smooth away from the diagonal, satisfying*

$$\varphi(p) = \int_M G_M(p, q) \Delta \varphi(q) dA(q) - \int_{\partial M} \varphi(q) \nabla_q G_M(p, q) \cdot \mathbf{n} d\ell(q)$$

for every $\varphi \in C^2(\bar{M})$. Furthermore, for every $p \in \bar{M}$ there exists a constant $c > 0$ such that $|G_M(p, q)| < c(1 + |\log d(p, q)|)$.

For $\mu \in \mathcal{M}(M)$ the function $u := G_M * \mu$ is a distributional solution to $\Delta u = \mu$ with zero Dirichlet boundary conditions. Hence, we can define a *solution operator* $S\mu := G_M * \mu$. Moreover, the following classical result shows that S is a bounded linear operator from $\mathcal{M}(M)$ into $W_0^{1,p}(M)$.

THEOREM A.2 (LITTMAN ET AL. [1963, THEOREM 5.1]). *Let $\Omega \subset \mathbb{R}^2$ be a bounded open set with $\partial\Omega$ of class C^∞ . Let μ be a signed Radon measure with finite total variation in Ω . Let L be a uniformly elliptic operator in divergence form. Then the Dirichlet problem*

$$Lu = \mu$$

with $u \in W_0^{1,1}(\Omega)$ has a unique solution with the property that $u \in W_0^{1,p}(\Omega)$ with $\|\nabla u\|_{L^p(\Omega)} \leq c\|\mu\|_{\mathcal{M}}$ for all $1 < p < 2$.

By the Sobolev embedding theorem $W_0^{1,p}(M) \subset L^2(M)$, and so we will actually consider $S : \mathcal{M}(M) \rightarrow L^2(M)$. By localizing, we deduce that the same result holds for the solution operator associated to the Laplace-Beltrami operator on (M, g) . The result is shown using the method of duality introduced by Stampacchia [Littman et al. 1963; Stampacchia 1965]. In particular, S is constructed as the Banach space adjoint of the (bounded linear) operator $S_* : W^{-1,p'}(M) \rightarrow C(M)$. Again, by using the chain of embeddings and isomorphisms $L^2(M) \cong (L^2(M))^* \hookrightarrow W^{-1,p'}(M)$ we will actually consider S_* as a map from $L^2(M)$ into $C(M)$. Since Δ is self-adjoint the map S_* is also the solution operator to the Poisson equation $\Delta u = \sigma$ (again, with zero Dirichlet boundary conditions), where now $\sigma \in L^2(M)$ —the fact that S is the adjoint $(S_*)^*$ is important for our application of Fenchel-Rockafellar duality.

A.1.1 Boundary Conditions. There are many practical benefits to jointly optimizing over both boundary conditions and cone configurations. Rigorously treating the boundary conditions in the continuous setting is quite technical; nevertheless, the treatment of boundary conditions numerically is extremely simple—see Sec. 5.3.4 for the necessary practical details.

We extend the solution operator to allow more general Dirichlet boundary data $b \in H^{3/2}(M)$. Since $b \in H^{3/2}(M)$ we can employ usual functional analytic methods (Lax-Milgram) to deduce the existence of a weak solution to

$$\begin{aligned} \Delta v &= 0 && \text{in } M, \\ v &= b && \text{on } \partial M. \end{aligned}$$

The boundary conditions above are understood in the trace sense. Furthermore, elliptic theory and the Sobolev trace theorem gives us the bound $\|v\|_{H^2(M)} \leq c\|b\|_{H^{3/2}(\partial M)}$. Thus, we have a continuous solution operator $B : H^{3/2}(\partial M) \rightarrow H^2(M)$ to the above boundary value problem. We choose to only consider boundary data in $H^{3/2}(\partial M)$ since solutions to the above boundary value problem are in $H^2(M) \hookrightarrow C(M)$, which considerably simplifies the boundary conditions in our pre-dual problem.

Combining the solution operators S and B we can construct solutions to the Poisson equation

$$\begin{aligned} \Delta u &= \mu && \text{in } M, \\ u &= b && \text{on } \partial M \end{aligned}$$

by setting $u = S\mu + Bb$. We will denote this solution operator $S_\partial : \mathcal{M}(M) \times H^{3/2}(\partial M) \rightarrow X$. We write X for the Hilbert space $L^2(M) \oplus H^{3/2}(\partial M)$ (which we identify with its dual X^* through the Riesz representation theorem).

A.2 Generalizations

Before constructing the pre-dual problem and computing the convex conjugates, we discuss the mathematical framework for treating the generalizations that appear in Sec. 7.

Optimizing the Boundary. Using the generalized solution operator above we can properly include optimizations over Dirichlet boundary data in $H^{3/2}(\partial M)$. Alternatively, we can optimize over Neumann boundary data—in light of the Cherrier boundary conditions (Sec. 3.1), penalizing the measure norm on the boundary data promotes polygonal boundaries in the optimization. In this case, one can only (formally) specify measure data on the boundary. A more careful treatment of measure valued Neumann boundary conditions is still necessary for a proper analysis of the optimization problem in the continuous setting. From a numerical perspective, this poses no challenges and we use this to penalize sparsity of the geodesic curvature on ∂M .

Visibility. Consider a closed subset $U_R \subset M$, and let $\mathcal{M}(U_R) \subset \mathcal{M}(M)$ denote the set of measures with support in U_R . This change enforces geometric constraints on the locations of the optimal cones.

Similarly, consider an open subset $U_E \subset M$. We can optimize over cone configurations that minimize the area distortion only in U_E by considering the energy $\tilde{E}(u) := E(\chi_{U_E} \cdot u)$.

M	space of measures
E	area distortion energy
R	measure space regularizer
λ	tuning parameter
U_R	region where cones can be placed
U_E	region where distortion is measured
S	solution operator to $\Delta u = \mu$
X	$L^2(M) \oplus H^{3/2}(\partial M)$
G_M	Green potential
s_i	free cone singularity sites
K	discrete triangle mesh
S_K	FEM semidiscretization of S
$P_{w_{\mathcal{R}}}$	projection operator (Eqn. 18)

Fig. A.1. The notation we use throughout this paper is summarized here.

Prescribed Singularities. To freely allow cone singularities at a finite number of sites $\{s_1, \dots, s_k\} \subset M$, while optimizing over additional cones, we replace the measure-norm by the regularization

$$R(\mu) := |\mu|(U_R \setminus \{s_1, \dots, s_k\}).$$

Thus, there is no cost to placing cone singularities at these sites. $|\mu|$ denotes the total variation measure of μ .

We introduce positive weight functions $w_{\mathcal{E}}$ and $w_{\mathcal{R}}$ (bounded away from zero) to add soft geometric constraints on the optimal cones. Now, we define

$$E(u) := \frac{1}{2} \int_M |u|^2 w_{\mathcal{E}} dA \quad \text{and} \quad R(\mu) := \int_M w_{\mathcal{R}} d|\mu|,$$

which are the weighted area distortion measure and weighted regularization, respectively. Under these modifications, the optimal control problem prefers to place cone singularities where $w_{\mathcal{R}}$ is small in a way that reduces the area distortion where $w_{\mathcal{E}}$ is large. See Sec. 8.1 for a practical geometric example.

A.3 Applying Convex Duality

In Sec. 4.3 we claim that the Fenchel-Rockafellar dual of Problem 10 is Problem 9. We carry out the computations for the convex conjugates and show how to realize the optimality conditions as weak solutions to PDEs. For notational simplicity we set $U_E = U_R = M$ and consider only zero Dirichlet boundary conditions, but otherwise consider the weights from the generalized problem presented in App. A.2. Adding in these additional generalizations follows in exactly the same way.

Recall the following relationship between the convex subdifferential and the convex conjugate.

PROPOSITION A.3 (FENCHEL-YOUNG INEQUALITY). *Consider a proper convex function $F : X \rightarrow \overline{\mathbb{R}}$.*

$$F^*(L) + F(x) \geq L(x) \quad \text{for all } x \in X, L \in X^*.$$

Equality holds if and only if $L \in \partial F(x)$.

Convex Conjugates (Area Distortion). Since the area distortion energy is defined on the Hilbert space $L^2(M)$ we can simply compute the convex conjugate of E to obtain the corresponding term in the

pre-dual problem. The Fenchel-Moreau theorem states that $E^{**} = E$, and so this justifies this procedure. It is straightforward to see that E is Fréchet differentiable with

$$dE(u) = uw_{\mathcal{E}},$$

where we identify $L^2(M) \cong (L^2(M))^*$ in the usual way. Thus, by Proposition A.3 the convex conjugate of E is given by

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

Convex Conjugates (Measure Space Regularizer). To determine the appropriate pre-dual for the measure space regularizer we considered the simple finite-dimensional case where the dual of box-constraints dualize to L^1 -norm. Recall that since we need to make the change of variables $\sigma := K dA - \mu$ (see Sec. 4.3) we need to construct G such that $G^*(\sigma) = R(K dA - \sigma)$.

Thus, define the convex functional $G : C(M) \rightarrow \overline{\mathbb{R}}$ by

$$G(\varphi) := \mathbb{1}_{\lambda w_{\mathcal{R}}}(\varphi) - \int_M K \varphi dA,$$

where analogous to $\mathbb{1}_{\lambda}$ we consider the indicator function

$$\mathbb{1}_{\lambda w_{\mathcal{R}}}(\varphi) := \begin{cases} 0 & \text{if } |\varphi(\mathbf{x})| \leq \lambda w_{\mathcal{R}}(\mathbf{x}) \text{ for all } \mathbf{x} \in M, \\ +\infty & \text{else.} \end{cases}$$

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

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

Let M^+ denote the support of $(\mu + \Omega)^+$ and M^- denote the support of $(\mu + \Omega)^-$. The Jordan decomposition theorem states that these sets are disjoint. Construct a sequence of continuous functions $\{\varphi_n\}_{n=1}^{\infty} \subset C(M)$ such that pointwise

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

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

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

The reverse inequality is immediate from the inequality constraints on φ . Summarizing, we have shown that $G^*(\mu) = \lambda R(K dA + \mu)$ is the desired measure space regularization appearing in our relaxed formulation.

Applying Fenchel-Rockafellar duality. We will use the duality theorem in the form presented in Sec. 3.4—using the same notation we will take $F = E^*$ and G as above. Furthermore, we will take $\Lambda^* = S : \mathcal{M}(M) \rightarrow L^2(M)$ to be the associated solution operator. Now the primal problem from the duality theorem reads

$$\inf_{u \in L^2(M)} E^*(u) + G(\Lambda u).$$

Set $\varphi = \Lambda u$, and note that φ is a weak solution to

$$\begin{aligned} \Delta \varphi &= u & \text{in } M, \\ \varphi &= 0 & \text{on } \partial M, \end{aligned}$$

where the boundary conditions are understood in the trace sense. Using this formulation of Λ , we can expand out this form of the primal problem to obtain

$$\begin{aligned} & \underset{u \in L^2(M)}{\text{minimize}} && \frac{1}{2} \int_M \frac{1}{w_{\mathcal{E}}} |u|^2 dA - \int_M (K\varphi) w_{\mathcal{R}} dA \\ & \text{subject to} && \Delta\varphi = u && \text{in } M, \\ & && \varphi = 0 && \text{on } \partial M, \\ & && |\varphi(\mathbf{x})| \leq \lambda w_{\mathcal{R}}(\mathbf{x}) && \text{for all } \mathbf{x} \in M. \end{aligned} \quad (\text{A.1})$$

The necessary technical conditions that there exists some $u_0 \in X$ such that $E^*(u_0) + G(\Lambda u_0) < +\infty$ with G continuous at Λu_0 holds trivially for $u_0 = 0$. Applying the Fenchel-Rockafellar duality theorem yields the equivalent optimization

$$- \min_{\sigma \in \mathcal{M}(M)} E(\Lambda^* \sigma) + G^*(-\sigma).$$

Again, we expand this out using our computations of the convex conjugates and the definition of the solution operator S to find

$$\begin{aligned} & \underset{\sigma \in \mathcal{M}(M)}{\text{minimize}} && \frac{1}{2} \int_M |u|^2 w_{\mathcal{E}} dA + \lambda \int_M w_{\mathcal{R}} d|K dA - \sigma| \\ & \text{subject to} && \Delta u = \sigma && \text{in } M, \\ & && u = 0 && \text{on } \partial M, \end{aligned} \quad (\text{A.2})$$

Optimality Conditions. Finally, we derive the optimality conditions that relate the minimizers of the optimization problems Problem A.1 and Problem A.2. The Fenchel-Rockafellar duality theorem provides the first order necessary and sufficient optimality conditions

$$\begin{cases} \Lambda^* \bar{\sigma} & \in & \partial E^*(\bar{u}), \\ -\bar{\sigma} & \in & \partial G(\Lambda \bar{u}). \end{cases}$$

Setting $\bar{\varphi} = \Lambda \bar{u}$, using the definition of the solution operators Λ and Λ^* , and using the fact that $dE^*(v) = v$ we can rewrite the optimality conditions as

$$\begin{cases} \Delta \bar{u} = \bar{\sigma} & \text{with } \bar{u}|_{\partial M} = 0, \\ \Delta \bar{\varphi} = w_{\mathcal{E}} \bar{u} & \text{with } \bar{\varphi}|_{\partial M} = 0, \\ K dA - \bar{\sigma} \in \partial \mathbb{1}_{\lambda w_{\mathcal{R}}}(\bar{\varphi}). \end{cases}$$

Substituting $\mu = K dA - \sigma$ into the above provides the optimality system between the primal problem and the relaxed optimization formulated with the singular Yamabe equation as the PDE constraint:

$$\begin{cases} \Delta \bar{u} = K dA - \bar{\mu} & \text{with } \bar{u}|_{\partial M} = 0, \\ \Delta \bar{\varphi} = w_{\mathcal{E}} \bar{u} & \text{with } \bar{\varphi}|_{\partial M} = 0, \\ \bar{\mu} \in \partial \mathbb{1}_{\lambda w_{\mathcal{R}}}(\bar{\varphi}). \end{cases}$$

Discretizing and solving this optimality system directly will be the basis of our algorithm for computing optimal cones.

Generalized Optimality Conditions. Working through the computations of the Fenchel-Rockafellar duality theorem with all of the generalizations from App. A.2 yields the following optimality system

$$\begin{cases} \Delta \bar{u} = K dA - \bar{\mu} & \text{with } \bar{u}|_{\partial M} = b, \\ \Delta \bar{\varphi} = w_{\mathcal{E}} \chi_{U_E} \bar{u} & \text{with } \bar{\varphi}|_{\partial M} = b, \\ \bar{\mu} \in \partial \mathbb{1}_{\lambda w_{\mathcal{R}}}(\chi_{U_R} \bar{\varphi}). \end{cases}$$

Here $b \in H^{3/2}(\partial M)$ are new optimization variables representing the boundary conditions, and χ_A denotes the characteristic function of the region A (i.e., $\chi_A(x)$ is 0 if $x \notin A$ and 1 if $x \in A$).

B DISCRETIZATION (DETAILS)

Recall that our algorithm takes in a triangle mesh $K = (V, E, F)$. Properly discretizing the optimality conditions is quite subtle due to the presence of measures. In this appendix we will discuss both how to derive the discretized optimality system as well as the detailed computations needed to reformulate the optimality conditions as a semismooth operator equation (see Sec. 5.3).

B.1 Discretized Optimality Conditions

To recover the cone configurations, we will need an appropriate discretization of the space of measures—in particular, we need to discretize the optimality conditions in a way that captures the central features of the optimization over the space of measures. For simplicity of presentation, we present the derivation only in the basic formulation of our problem. We utilize the conforming discretization of $\mathcal{M}(M)$ introduced by Casas, Clason, and Kunisch [2012]. We recall the basic framework here. Let Y_K denote the finite dimensional subspace of $L^2(M)$ given by linear combinations of the piecewise linear Whitney hat functions centered at the vertices V . Motivated by the Riesz representation theorem for $C(M)^*$, the space of measures is discretized as $Z_K := Y_K^*$ where the duality pairing is given by the duality pairing between $\mathcal{M}(M)$ and $C(M)$. Explicitly,

$$Z_K := \left\{ \sum_{i \in V} \theta_i \delta_i : \theta_i \in \mathbb{R} \right\}.$$

We now utilize this approach with the *variational discretization* technique of Hinze [2005]. The main idea is to *not* discretize the space of measures initially, but instead to only use the above discretization in the end to recover the cone configurations. We begin by considering the semidiscrete finite-element approximation of the solution operator—the term semidiscrete is used to emphasize that the space of measures is not discretized, whereas the state space $L^2(M)$ is. For a given measure $\mu \in \mathcal{M}(M)$, the discrete log-conformal factors $u = S_K \mu \in Y_K$ are characterized by

$$(\nabla u, \nabla \varphi)_{L^2(M; TM)} = \int_M \varphi d\mu$$

for all $\varphi \in Y_K$. So the semidiscrete solution operator is a map $S_K : \mathcal{M}(M) \rightarrow Y_K$. Now we consider the semidiscrete optimization problem

$$\underset{\mu \in \mathcal{M}(M)}{\text{minimize}} E(S_K \mu) + \lambda \|K dA - \mu\|_{\mathcal{M}}. \quad (\text{B.1})$$

Minimizers in $\mathcal{M}(M)$ exist by a standard application of the direct method of the calculus of variations. Notice, though, that S_K maps an infinite dimensional vector space into a finite-dimensional space. Thus, minimizers of Problem B.1 are not unique. However, Casas et al. [2012, Theorem 3.2] show that there is a unique $\bar{\mu} \in \mathcal{M}(M)$ satisfying $\bar{u} = S_K \bar{\mu}$ that is a linear combination of Dirac delta measures concentrated on the vertices in V —this provides a numerical accessible subspace of $\mathcal{M}(M)$ to look for minimizers.

Crucially note that this approach is fundamentally different from naively treating measures as integrable functions and utilizing standard finite-elements to obtain sparse solutions. At the level of both the semidiscrete optimization and the fully discrete optimality conditions the continuous theory regarding the space of measures is

used to determine discrete minimizers that faithfully represent minimizers in the smooth setting.

By restricting the search space to Z_K we can express solutions to Problem B.1 in terms of the coefficients of $\bar{\mu}$ and \bar{u} . By applying Fenchel-Rockafellar duality to the semidiscrete optimization problem (exactly as in Sec. 4.3 and App. A.3) yields the discrete optimality conditions

$$\begin{cases} \mathbf{L}u = \Omega - \mu, \\ \mathbf{L}^\top \varphi = \mathbf{M}u, \\ \mu \in \partial \mathbb{1}_\lambda(\varphi), \end{cases} \quad (\text{B.2})$$

where \mathbf{L} is the stiffness matrix of the Laplace-Beltrami operator (*i.e.*, the cotan-Laplacian), \mathbf{M} is the Galerkin mass matrix, and Ω is the discrete curvature 2-form.

We now describe our discretization of the sequence of regularized problems presented in Sec. 4.6. In a manner completely analogous to the discretization above, the starting point is to construct a semidiscrete regularized problem of Problem 9. Again, minimizers of the semidiscrete optimization problem can be guaranteed using standard variational tools. For this sequence of regularized problems to be useful, we need our discretization of the semidiscrete problems to converge to Problem B.1 as $\gamma \rightarrow 0^+$. Thus, once we have *any* minimizer, we store the values at the vertices integrated over the associated dual cell. By utilizing Fenchel-Rockafellar duality (as in the rest of the paper) we obtain that the fully discrete optimality system reads

$$\begin{cases} \mathbf{L}u = \Omega - \mu, \\ \mathbf{L}^\top \varphi = \mathbf{M}u, \\ \varphi - \gamma \mu \in \partial R(\mu), \end{cases} \quad (\text{B.3})$$

where here R is given by $R(x) = \|\mathbf{w}_{\mathcal{R}} x\|_{\ell^1}$, where the multiplication is understood pointwise, *i.e.*,

$$(\mathbf{w}_{\mathcal{R}} x)_i = (\mathbf{w}_{\mathcal{R}})_i x_i.$$

B.2 Semismooth Reformulation

As mentioned in Sec. 5.3 our algorithm will be based on applying a semismooth Newton method to obtain solutions that satisfy the discrete optimality systems System B.2 and System B.3.

The main challenge is in discretizing the subdifferential relationships. Since these optimality systems are framed over the space $\mathbb{R}^{|\mathcal{V}|}$ we can utilize *Moreau's proximal map* to reformulate these relationships as equality relationships. Our primary reference regarding the proximal map is Bauschke and Combettes [2017]. Throughout this section H is a Hilbert space.

Definition B.1. Let $f : H \rightarrow H$ be a convex function. *Moreau's 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 write $\text{Prox}_\gamma := \text{Prox}_{f/\gamma}$ when f is understood from context.

The main connection between the proximal map and subdifferential relationships is given by the following important relationship.

PROPOSITION B.2 ([BAUSCHKE AND COMBETTES 2017, PROPOSITION 23.2, EXAMPLE 23.3]). *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,*

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

The Regularized Problems. We begin by reformulating the optimality conditions for the regularized problem. By comparing this relationship with the subdifferential relationship in System B.3 we see that this is *exactly* the form we desire. From Proposition B.2 we see that

$$\varphi - \gamma \mu = \gamma \left(\frac{1}{\gamma} \varphi - \mu \right) \in \partial R(\mu)$$

is equivalent to

$$\mu = \text{Prox}_\gamma \left(\frac{\varphi}{\gamma} \right).$$

Now, we are only left with computing the proximal map of the discrete regularization $R : \mathbb{R}^{|\mathcal{V}|} \rightarrow \mathbb{R}^{|\mathcal{V}|}$.

PROPOSITION B.3.

$$\text{Prox}_{R/\gamma}(x) = \max \left(0, x - \frac{1}{\gamma} \mathbf{w}_{\mathcal{R}} \right) + \min \left(0, x + \frac{1}{\gamma} \mathbf{w}_{\mathcal{R}} \right).$$

PROOF. To find $\text{Prox}_\gamma(x)$ we need to minimize the functional

$$y \mapsto \sum_{i \in \mathcal{V}} \frac{\gamma}{2} |y_i - x_i|^2 + (\mathbf{w}_{\mathcal{R}})_i |y_i|.$$

It suffices to minimize the summand for every $i \in \mathcal{V}$. That is, we minimize the one-dimensional functional $y \mapsto \frac{\gamma}{2} (y - x_i)^2 + (\mathbf{w}_{\mathcal{R}})_i |y|$, which gives rise to

$$y = \begin{cases} 0 & \text{if } |x_i| \leq \frac{1}{\gamma} (\mathbf{w}_{\mathcal{R}})_i, \\ x_i - \frac{1}{\gamma} \text{sgn}(x_i) (\mathbf{w}_{\mathcal{R}})_i & \text{else.} \end{cases}$$

So we obtain the claimed formula for the proximal map, where the max and min are understood pointwise. \square

Using this, we obtain that

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

So we can reformulate System B.3 as

$$\begin{cases} \mathbf{L}u = \Omega - \mu, \\ \mathbf{L}^\top \varphi = \mathbf{M}u, \\ \gamma \mu = \max(0, x - \mathbf{w}_{\mathcal{R}}) + \min(0, x + \mathbf{w}_{\mathcal{R}}) \end{cases} \quad (\text{B.4})$$

To simplify notation, we consider $P_{\mathbf{w}_{\mathcal{R}}}(\varphi) := \max(0, \varphi - \mathbf{w}_{\mathcal{R}}) + \min(0, \varphi + \mathbf{w}_{\mathcal{R}})$.

Now, System B.4 can be concisely encoded in the operator equation $F_\gamma(u, \varphi, \mu) = 0$ where F_γ is given by

$$F_\gamma(u, \varphi, \mu) = \begin{bmatrix} \mathbf{L}u - \Omega + \mu \\ \mathbf{L}^\top \varphi - \mathbf{M}u \\ \gamma \mu - P_{\mathbf{w}_{\mathcal{R}}}(\varphi) \end{bmatrix}.$$

In fact, we can simplify this even further by substituting the third equality into the first. That is, we can consider

$$F_\gamma(u, \varphi) = \begin{bmatrix} \mathbf{L}u - \Omega + \frac{1}{\gamma} P_\lambda(\varphi) \\ \mathbf{L}^\top \varphi - \mathbf{M}u \end{bmatrix}.$$

The Relaxed and Unregularized Problem. As above, we want to reformulate System B.2 as an operator equation of the form $F(u, \varphi, \mu) = 0$. Again, we will utilize Proposition B.2 to help reformulate

$$\mu \in \partial \mathbb{1}_{w_{\mathcal{R}}}(\varphi).$$

By Proposition B.2, for any $\gamma > 0$

$$\varphi = \text{Prox}_\gamma \left(\varphi + \frac{1}{\gamma} \mu \right)$$

is equivalent to

$$\gamma \left(\varphi + \frac{1}{\gamma} \mu - \varphi \right) = \mu \in \partial \mathbb{1}_{w_{\mathcal{R}}}(\varphi).$$

Thus, we can write System B.2 as

$$\begin{cases} \mathbf{L}u = \Omega - \mu, \\ \mathbf{L}^\top \varphi = \mathbf{M}u, \\ \varphi = \text{Prox}_\gamma \left(\varphi + \frac{1}{\gamma} \mu \right), \end{cases} \quad (\text{B.5})$$

for any $\gamma > 0$. As before, we need to compute the proximal map for $\mathbb{1}_{w_{\mathcal{R}}}$. However, this is quite simple since this is just the indicator function of a convex set:

$$\text{Prox}_{\mathbb{1}_{w_{\mathcal{R}}}}(\varphi) = \arg \min_x \left(\frac{1}{2} \|x - \varphi\|^2 + \mathbb{1}_{w_{\mathcal{R}}}(x) \right),$$

and so we see that we are simply obtain the metric projection onto the box constraints defined by $w_{\mathcal{R}}$. The projection onto these box constraints can be easily described using the map $P_{w_{\mathcal{R}}}$:

$$\text{Proj}_{w_{\mathcal{R}}}(\varphi) = \varphi - P_{w_{\mathcal{R}}}(\varphi).$$

Thus, we can write

$$\varphi = \text{Prox}_\gamma \left(\varphi + \frac{1}{\gamma} \mu \right) = \varphi + \frac{1}{\gamma} \mu - P_{w_{\mathcal{R}}} \left(\varphi + \frac{1}{\gamma} \mu \right).$$

Since the above equivalence holds for any $\gamma > 0$, we can take $\gamma = 1$. Now, rearranging terms yields

$$\mu = P_{w_{\mathcal{R}}}(\varphi + \mu),$$

which allows us to express the optimality system as $F(u, \varphi, \mu) = 0$ where

$$F(u, \varphi, \mu) := \begin{bmatrix} \mathbf{L}u - \Omega - \mu \\ \mathbf{L}^\top \varphi - \mathbf{M}u \\ \mu - P_{w_{\mathcal{R}}}(\varphi + \mu) \end{bmatrix}.$$

C REGULARITY AND ROUNDING

In this section we prove a negative result regarding the sparsity structure of the minimizers—in particular, the minimizing measures will *never* represent cone singularities. We go on to rectify the situation by proving a stability result which justifies rounding arbitrary measures to nearby cone singularity configurations. On almost all of the examples in this paper the minimizing measure was already supported on isolated vertices.

We begin with a technical result, which relates the support of the minimizing measure to the values of the adjoint state.

PROPOSITION C.1. *Let $\bar{\mu}$ be the optimizing measure, and let $\bar{\varphi}$ be the optimal adjoint state. We have that*

$$\text{supp } \bar{\mu}^+ \subseteq \{ \mathbf{x} \in M : \bar{\varphi}(\mathbf{x}) = +\lambda w_{\mathcal{R}}(\mathbf{x}) \},$$

$$\text{supp } \bar{\mu}^- \subseteq \{ \mathbf{x} \in M : \bar{\varphi}(\mathbf{x}) = -\lambda w_{\mathcal{R}}(\mathbf{x}) \}.$$

PROOF. This result follows from the optimality conditions presented in System 11. In particular, we will use the condition that $\bar{\mu} \in \partial \mathbb{1}_{\lambda w_{\mathcal{R}}}(\bar{\varphi})$. From Proposition A.3 and this subdifferential relationship we have that

$$\langle \bar{\mu}, \bar{\varphi} \rangle = \mathbb{1}_{\lambda w_{\mathcal{R}}}(\bar{\varphi}) + \mathbb{1}_{\lambda w_{\mathcal{R}}}^*(\bar{\mu}) = \lambda \int_M w_{\mathcal{R}} d|\bar{\mu}|. \quad (\text{C.1})$$

In the second equality we use that fact that since $\mathbb{1}_{\lambda w_{\mathcal{R}}}$ is subdifferentiable at $\bar{\varphi}$ that $\mathbb{1}_{\lambda w_{\mathcal{R}}}(\bar{\varphi}) = 0$. 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. C.1 we obtain

$$\int_M (\text{sgn } \bar{\mu}) \bar{\varphi} d|\bar{\mu}| = \lambda \int_{U_R} w_{\mathcal{R}} d|\bar{\mu}|.$$

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

$$\text{supp } \bar{\mu}^+ \subseteq \{ \mathbf{x} \in U_R : \text{sgn } \bar{\mu} = +1 \} \subseteq \{ \mathbf{x} \in M : \bar{\varphi}(\mathbf{x}) = +\lambda w_{\mathcal{R}}(\mathbf{x}) \},$$

$$\text{supp } \bar{\mu}^- \subseteq \{ \mathbf{x} \in U_R : \text{sgn } \bar{\mu} = -1 \} \subseteq \{ \mathbf{x} \in M : \bar{\varphi}(\mathbf{x}) = -\lambda w_{\mathcal{R}}(\mathbf{x}) \},$$

□

We make the assumption that $U_R = U_E = M$. For simplicity we assume that the boundary conditions are fixed to $u = 0$ and that $w_{\mathcal{R}} \equiv 1$. We crucially assume that M is homeomorphic to the unit disk in \mathbb{R}^2 so we can perform an initial flattening—let u_0 be the log conformal factors which make $(M, e^{2u_0}g)$ flat. Using this initial flattening, we consider the local picture mentioned in Sec. 4.1.1.

The following lemma 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 the optimizing measure cannot contain Dirac deltas.

LEMMA C.2. *Let $\bar{\mu}$ be the optimizing measure of Problem 9, and let $\Delta \bar{u} = \bar{\mu}$ with $\bar{u} = 0$ on ∂M . Then $\bar{u} \in L^\infty(M)$ with*

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

PROOF. As before, consider the adjoint state $\bar{\varphi} \in C(M)$ that satisfies $\Delta_{\mathbb{R}^2} \bar{\varphi} = \bar{u} - u_0$ with $\bar{\varphi} = 0$ on ∂M . For notational convenience, set

$$B := \|u_0\|_{L^\infty(M)} + \left(\inf_{\mathbf{x} \in M} w_{\mathcal{G}}(\mathbf{x}) e^{2u_0(\mathbf{x})} \right)^{-1}$$

Assume, for the sake of contradiction, $\|\bar{u}\|_{L^\infty(M)} > B$. Thus, there exists $\varepsilon > 0$ such that the set

$$A := \{ \mathbf{x} \in M : |\bar{u}(\mathbf{x})| > B + \varepsilon \}$$

has positive measure. Without loss of generality, we can assume that $\sup_{\mathbf{x} \in M} \bar{u}(\mathbf{x}) > B + \varepsilon$. By the maximum principle for logarithmic potentials [Saff and Totik 1997, Corollary 3.3] there exists some $\mathbf{x}_0 \in \text{supp } \bar{\mu}^+$ satisfying $\bar{u}(\mathbf{x}_0) > B + \varepsilon$. Now let $\delta > 0$ be such that $\text{dist}(\text{supp } \bar{\mu}^+, \text{supp } \bar{\mu}^-) > \delta$ (which follows from Proposition C.1). 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 $y \in C^\infty(M)$ 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

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

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 $\|\bar{u}\|_{L^\infty(M)} \leq B$, as desired. \square

Using the above L^∞ -bound on the optimal log-conformal factors we will be able to bound the H^1 -norm of u . We begin by bounding the H^1 -norm of u when the measure μ is absolutely continuous with respect to the area element dA :

LEMMA C.3. *Let $f \in C_c^\infty(M; \mathbb{R})$, and consider the measure μ given by $d\mu = f \, d\mathbf{x}$. Then*

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

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

PROOF. Since f is smooth, we know that u is smooth; furthermore, u vanishes on ∂M . Integration by parts reveals

$$\begin{aligned} \int_M u(\mathbf{x}) \, d\mu(\mathbf{x}) &= \int_M u(\mathbf{x}) f(\mathbf{x}) \, d\mathbf{x} \\ &= \int_M u(\mathbf{x}) \Delta u(\mathbf{x}) \, d\mathbf{x} = \int_M \|\nabla u(\mathbf{x})\|^2 \, d\mathbf{x}. \end{aligned}$$

\square

By an approximation argument we can apply the previous result to obtain the desired regularity result:

THEOREM C.4. *The optimizing measure $\bar{\mu} \in H_0^1(M)^*$.*

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

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

Note that the right hand side is finite by Lemma C.2. Note that by Lemma C.3 and a simple approximation argument (approximate $\bar{\mu}$ by smooth functions, and pass to the limit) we obtain

$$\|\nabla \bar{u}\|_{L^2(M; \mathbb{R}^2)}^2 \leq \int_M u(\mathbf{y}) \, d\bar{\mu}(\mathbf{y}),$$

which shows that $\bar{u} \in H_0^1(M)$. Since $\Delta : H_0^1(M) \rightarrow H_0^1(M)^*$ we conclude that $\bar{\mu} \in H_0^1(M)^*$. \square

This technical result is important since it tells us that the optimizing measure cannot consist of *any* Dirac deltas. 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.

Rounding. Although in the continuous setting the minimizing measures never consist of Dirac delta measures, we find that in practice they often are delta measures. To resolve these conflicting perspectives, we show that we can round the minimizing measure to a collection of delta measures in a way that does not significantly reduce the area distortion. Rigorously, we utilize the framework of optimal transportation.

Recall that the Wasserstein distance between two probability measures μ, ν is given by

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

where d_g is the geodesic distance in M . On the other hand, we have that

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

where $\Delta u = \mu$ in M with $u|_{\partial M} \equiv 0$. Note that the H^{-2} norm is nothing more than the area distortion energy $E(u)$ from our cone optimization problem.

For any two functions f, g we use the notation $f \lesssim g$ to indicate that there is some universal constant c for which $f \leq cg$. Similarly, we write $f \sim g$ to mean that $f \lesssim g$ and $g \lesssim f$.

We first prove this 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 C.5. *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 $A_{\mathbf{y},z} := \{\mathbf{a} \in B(0, 2R) : |\mathbf{a}| \leq |\mathbf{a} + \mathbf{y} - \mathbf{z}|\}$. 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

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

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

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

COROLLARY C.6. *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. Follows immediately from Theorem C.5 and [Helms 2009, Theorem 3.2.12], which states that $G_Y(\mathbf{x}, \mathbf{y}) \leq G_{B(0, R)}(\mathbf{x}, \mathbf{y})$. \square

By working in local coordinates we also obtain the result for any compact manifold M .

Following [Ambrosio et al. 2011; Mainini 2012], 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.

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

LEMMA C.7. *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.*

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

COROLLARY C.8. *Let (M, g) be a compact 2-manifold. For every $\varepsilon > 0$ there exists a configuration of cone singularities such that the area distortion of the resulting conformal flattening less than ε .*

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

$$|E(u_\Omega) - E(u_\mu)| < \varepsilon$$

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 C.5. By Lemma C.7 we can find a sequence of linear combinations of Dirac delta measures $\{\mu_n\}_{n=1}^\infty$ such that $\mu_n \xrightarrow{*} \Omega$. By [Mainini 2012, Proposition 3.8] we deduce that

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

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

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

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

Proposition 5 in Myles and Zorin [2013] proves a stronger result, namely that there exists a *seamless* cone parameterization that has area distortion less than ε for every $\varepsilon > 0$.