

On the Convexity and Feasibility of the Bounded Distortion Harmonic Mapping Problem

Zohar Levi
Victoria University of Wellington, New Zealand

Ofir Weber
Bar Ilan University, Israel




Figure 1: *Cobra.* (a) source domain. (b) Input mapping generated with As-Rigid-As-Possible shape deformation [Sorkine and Alexa 2007], to be projected on the bounded distortion space with singular values bounded by $\sigma = 0.2, \Sigma = 5$. (d,e) The projection methods [Aigerman and Lipman 2013; Kovalsky et al. 2015] produce a mapping that is not locally injective (though the triangles are orientation preserving). (c,f) The projection methods based on [Lipman 2012; Chen and Weber 2015] failed to capture the correct orientation even after many iterations.

Abstract

Computation of mappings is a central building block in many geometry processing and graphics applications. The pursuit to compute mappings that are injective and have a controllable amount of conformal and isometric distortion is a long endeavor which has received significant attention by the scientific community in recent years. The difficulty of the problem stems from the fact that the space of bounded distortion mappings is nonconvex. In this paper, we consider the special case of harmonic mappings which have been used extensively in many graphics applications. We show that, somewhat surprisingly, the space of locally injective planar harmonic mappings with bounded conformal and isometric distortion has a convex characterization. We describe several projection operators that, given an arbitrary input mapping, are *guaranteed* to output a bounded distortion locally injective harmonic mapping that is closest to the input mapping in some special sense. In contrast to alternative approaches, the optimization problems that correspond to our projection operators are shown to be always feasible for any choice of distortion bounds. We use the boundary element method (BEM) to discretize the space of planar harmonic mappings and demonstrate the effectiveness of our approach through the application of planar shape deformation.

Keywords: injective mappings, conformal mappings, harmonic mappings, bounded distortion, shape deformation

Concepts: •Computing methodologies → Computer graphics; Shape modeling; Mesh models; Shape analysis;

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 ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. © 2016 ACM.

SIGGRAPH '16 Technical Paper, July 24–28, 2016, Anaheim, CA.

ISBN: 978-1-4503-4279-7/16/07

DOI: <http://dx.doi.org/10.1145/2897824.2925929>

1 Introduction

Computing a mapping encapsulates in it a myriad of geometry processing applications: shape deformation-and-interpolation, parameterization, and remeshing to name a few. Smooth mappings with low angular and metric distortion are often desired as they are visually more appealing. Furthermore, non-injective and degenerated mappings can jeopardize further geometric processing operations such as remeshing and physical simulations.

Most of the available methods do not provide guarantees on the injectivity of the mappings they produce nor can they assure that the distortion does not exceed a certain threshold, as the problem, in general, is considered hard. However, the trend in recent years has been to offer approximate solutions to this unsolved problem with a wide variety of properties [Lipman 2012; Weber et al. 2012; Aigerman and Lipman 2013; Aigerman et al. 2014; Poranne and Lipman 2014; Levi and Zorin 2014; Kovalsky et al. 2014; Chen and Weber 2015; Kovalsky et al. 2015]. A main insight that inspired these methods is that a problem can be solved efficiently if it is convex. More importantly, no initial feasible point is needed, and achieving a global minimum is guaranteed (if it exists). Unfortunately, the underlying bounded distortion problem leads to nonconvex constrained optimizations. To overcome this and take advantage of the convexity benefits, the main paradigm that is offered, for example, by [Lipman 2012] is to carve a maximal convex piece from the space of solutions, convexifying the problem. The convexification achieves its goal only partially. Many types of convex problems can be solved efficiently, but a choice of which convex piece of the original space to carve needs to be made, and the global minimum or even a feasible solution would likely not be in it.

Consider the following mock example to illustrate the concept of convexification: Find the closest point $z \in \mathbb{C}$ to a reference point z_0 such that z is inside an annulus (cyan domain in the inset) with an inner radius 1 and outer radius 2. More formally, we would like to solve the following optimization problem


$$\min_{z \in \mathbb{C}} |z - z_0| \quad (1a)$$

$$s.t. \quad 1 \leq |z| \quad (1b)$$

$$|z| \leq 2, \quad (1c)$$

While the energy (1a) is convex, the domain is not, hence the problem is nonconvex. A convexification of the problem in the spirit of [Lipman 2012] would substitute the nonconvex constraint (1b) with the half-plane (pink) $1 \leq \text{Re}(z)$. The new feasible domain (purple) is given by the intersection of the half-plane with the circle (1c) and it is a maximal convex subset of the annulus. The half-plane can be chosen to be tangent to the inner circle at any point.

The particular point is determined in [Lipman 2012] by a choice of a user parameter (so-called frame). The price for the convexification is apparent. We barely covered a quarter of the annulus. If for example we consider the reference point $z_0 = -0.5$, the global optimum would be obtained for $z = -1$, but the convexified optimization would return $z = 1$ instead. In fact, there is only one choice of half-plane that would yield the optimal solution.



Our insight in this paper is that by a change of variables, we can change the shape of the domain such that it becomes convex (Section 4) or nearly convex (Section 6). In our mock example, instead of representing z in Cartesian coordinates, consider the principal branch of the complex logarithm operator $\text{Log}(z) = \ln|z| + i\text{Arg}(z) = l + i\theta$ that maps the annulus *bijectively* to the convex rectangle $[0, \ln(2)] \times (-\pi, \pi]$. Expressing the energy (1a) in the new variables (l, θ) gives $|e^{l+i\theta} - z_0|$ which is not convex, however, we can design a different energy with similar geometric meaning that measures the Euclidean distance in the alternative rectangular domain. To this end, we also transform z_0 using the same Log operator. We now have the following *alternative* convex energy which is minimized over a convex domain

$$\min_{l, \theta} |l + i\theta - \text{Log}(-0.5)| \quad (2a)$$

$$s.t. \quad 0 \leq l \leq \ln(2) \quad (2b)$$

$$-\pi < \theta \leq \pi. \quad (2c)$$

The alternative rectangular domain is in one-to-one correspondence with the annulus, so no solutions are overlooked and the problem is feasible (as long as the original one is). The global minimizer is attained at $(l, \theta) = (0, \pi)$ which is then transformed back using the inverse operator, giving $z = e^{i\pi} = -1$.

2 Related Work

The computation of mappings is essential in many graphics applications such as shape deformation, shape interpolation, parameterization, remeshing, and quadrangulation. The amount of literature is vast and we do not aim at thoroughly reviewing it. We refer the reader to comprehensive surveys depending on the application in hand: Parameterization [Floater and Hormann 2005; Sheffer et al. 2006], deformation [Sorkine 2006; Botsch and Sorkine 2008], shape interpolation [Wolberg 1998; Alexa 2002] and quadrangulation [Bommes et al. 2013]. All these applications share a common ground. They all involve the computation of a mapping between two domains. They differ by the domains they consider, the constraints they employ, and the properties of the mappings they attempt to compute. We focus mainly on the works most relevant to ours, in the sense that they target very specific properties of the underlying mappings. These mappings are bijective (mostly locally) and they provide control over the induced amount of angular and/or metric distortion.

A popular approach to guarantee global bijectivity is based on a classic result by Tutte [1963] that embeds a disk-like mesh in a convex boundary shape. Floater [1997] generalized the method to

better capture the geometric properties of the mesh, but the obtained mappings tend to be highly distorted when a shape with strong boundary concavities is unnaturally forced to become convex. These kind of mappings can be viewed as discrete harmonic mappings, where the smooth analogue to Tutte’s result is the Radó theorem [Duren 2004]. [Weber and Zorin 2014] and [Aigerman et al. 2014] bijectively map both the source and the target domains to an intermediate convex domain using discrete harmonic mappings, then invert one mapping and compose with the other to obtain a bijection between nonconvex domains.

Conformal mappings are popular, since they are locally injective and have zero angular distortion. Linear methods such as [Lévy et al. 2002; Weber et al. 2009] cannot guarantee local injectivity while nonlinear (yet convex) methods such as [Springborn et al. 2008; Kharevych et al. 2006] typically achieve injectivity when the input mesh is of good quality. Yet, conformal mappings tend to introduce excessive scale variations and are somewhat too strict. The more general class of quasi-conformal mappings is often considered where a bounded amount of conformal distortion is allowed and is typically traded-off for a lower scale distortion. The mappings we compute belong to that class and additionally ensure that the amount of scale is bounded without compromising smoothness. Our method can be viewed as a generalization of [Weber and Gotsman 2010], who parameterized the space of conformal mappings to the broader class of quasi-conformal harmonic mappings.

The work of Weber et al. [2012] computes an optimal (extremal) quasi-conformal mapping given fixed boundary constraints and is quite robust, though the method is nonlinear and nonconvex. Other nonlinear approaches such as [Hormann and Greiner 2000; Schüller et al. 2013] attempt to obtain local injectivity by incorporating infinite penalty for inverted elements. The problem with this approach is that it requires a locally injective mapping to begin with. In contrast, our algorithm can take any input mapping and produce in return an injective one.

The main difficulty is that the underlying constraints for local injectivity or bounded distortion are inherently nonconvex and cannot be enforced efficiently in practice. Nonconvex optimization is considered a hard problem in computer science and mathematics. A popular and successful approach to address nonconvex problems of the concerned type is given by the work of Lipman [2012] for 2D manifolds and later extended to volumetric mappings in [Kovalsky et al. 2014]. Rather than solving the inaccessible nonconvex problem, a convex problem is considered. The convex problem is specifically tailored by substituting the nonconvex constraints with convex ones, such that they define a maximal convex subspace within the nonconvex space. Finding a point in this subspace becomes easy, and off-the-shelf convex optimization solvers can be employed. The convex subspace is chosen according to some user defined parameters (the so-called local frames), which dramatically affects the final outcome which can be one of two: i) the convex region is nonempty. In that case the problem is considered feasible and the global minimum (within the restricted convex subspace) is attained. ii) the convex region is empty. The former outcome is clearly preferable. Moreover, it allows for adaptation of the parameters, followed by additional convex solves which may lead to further decrease of the energy. The latter case is a complete failure, leaving us with nothing. Hence, to completely avoid it, the parameters are often chosen in a way that ensures feasibility. For planar mappings the only known sensible choice is the identity mapping [Poranne and Lipman 2014; Chen and Weber 2015] while for surface parameterization, a Tutte’s embedding has been used [Aigerman et al. 2014]. Such a strategy often fails to obtain an optimal result even when many iterations are performed. Another problem that arises is that the obtained solution typically lies on the boundary of the feasible domain, which often corresponds to a nonsmooth

mapping.

In contrast, our algorithms solve convex problems without relying on parameters and in a noniterative fashion, which is possible due to a clever choice of alternative spaces to operate within. These spaces are unique to harmonic mappings and are proven to be feasible for any specified bounds. Operating in these spaces has two additional benefits. First, any outcome is a C^∞ harmonic mapping, hence smoothness is not compromised even in tough cases. Second, the convex optimizations we solve are formulated as boundary value problems in a low dimensional space where similar to [Chen and Weber 2015], the inequality constraints are enforced only on the boundary of the domain, leading to highly-efficient algorithms.

Another approach for computing mappings with bounded distortion is by approximation. Aigerman and Lipman [2013] approximate the bounded distortion space at a vicinity of a given point with a convex space. The input to the method is a mapping with high distortion which is projected to the bounded distortion space by successively solving quadratic programs. Upon convergence, the result is guaranteed to be orientation-preserving bounded distortion, albeit convergence is not guaranteed.

[Kovalsky et al. 2015] uses a similar approach but employs a remarkable algebraic manipulation to accelerate the projection. In a sense, it combines the speed of alternating minimization techniques [Liu et al. 2008] with the low number of iterations usually required by the Gauss-Newton approach. The results are similar in nature to those of [Aigerman and Lipman 2013], but are computed more efficiently. Both methods are useful in a variety of situations and can handle challenging 2D and 3D inputs, but just like other nonlinear approaches, they cannot guarantee satisfaction of the constraints or optimality. Moreover, convergence of the algorithm guarantees that the elements (triangles or tets) are not inverted. Nevertheless, local injectivity is not guaranteed, since positive orientation is a necessary but not a sufficient condition for local injectivity at *vertices*. See the last paragraph of Section 3.3 in [Aigerman et al. 2014] for a discussion. Figure 2 shows a comparison of our method with four other methods that project an input of a holomorphic mapping with a singular point of vanishing derivative.

The above mentioned projection methods result in a non-locally injective mapping with the singular vertex left intact. The inset shows a zoom-in on the singular vertex of the result by Aigerman et al. Note that all the surrounding triangles are positively oriented but the vertex angle sum is 4π .

We conclude this section by referring to [Chien et al. 2016] which employs similar mathematical machinery and insights on harmonic planar mappings to perform bounded distortion shape interpolation.

3 Background

For completeness, this section provides an overview of known mathematical concepts and facts that will be used throughout the paper. For further reading we refer to the excellent books by Duren [2004] and Ahlfors [1979].

Harmonic and holomorphic functions. A real-valued function $u(x, y)$ is harmonic if it satisfies the Laplace equation

$$\Delta u = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} = 0. \quad (3)$$

A planar mapping is a vector function $[u(x, y), v(x, y)]$ from a region in the xy -plane to a region in the uv -plane. It is harmonic if both u and v satisfy (3). It would be convenient to use complex notations and to further rely on complex analysis. Let $z = x + iy$;

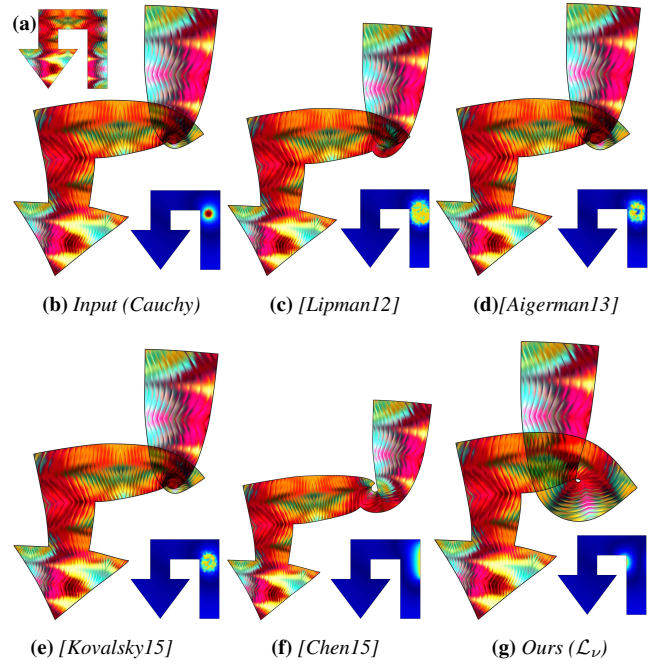


Figure 2: Arrow Twist. Isometric distortion visualized on the source domain. (a) Source domain. (b) non-injective input holomorphic mapping with a singularity, generated by [Weber et al. 2009]. (d,e) not locally injective. (c,f) locally injective but the tail of the arrow is wrongly rotated. [Lipman12] pushes the singularity to the boundary while [Chen15] prefers smoothness.

then we have $f(z) = u(z) + iv(z)$. Thus, a complex-valued harmonic function $f : \Omega \subset \mathbb{C} \rightarrow \mathbb{C}$ can be interpreted as a harmonic planar mapping. A complex-valued function $f = u + iv$ is holomorphic (also known as analytic) in a domain $\Omega \subset \mathbb{C}$ if it satisfies the Cauchy-Riemann equations at every point $z \in \Omega$

$$\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y}, \quad \frac{\partial u}{\partial y} = -\frac{\partial v}{\partial x}.$$

Holomorphic functions are the central objects of study in complex analysis. They are harmonic complex-valued functions (the converse is not true in general) and possess many strong and useful mathematical properties. They are differentiable and integrable (on simply-connected domains), and their derivative and anti-derivative are also holomorphic. The sums, products, and compositions of holomorphic functions are also holomorphic, and the quotient of holomorphic functions is holomorphic wherever the denominator does not vanish. Closure under complex conjugation does not hold. However, conjugating a holomorphic function results in an anti-holomorphic function which is closed under the same operations.

Derivatives. For an arbitrary differentiable complex function $f(z)$, the Wirtinger derivatives are defined as $f_z = \frac{1}{2}(f_x - if_y)$ and $f_{\bar{z}} = \frac{1}{2}(f_x + if_y)$. A direct consequence of the Cauchy-Riemann equations is that holomorphic functions satisfy $f_{\bar{z}}(z) = 0$, hence, it is convenient to denote the complex derivative of a holomorphic function by $f'(z) = f_z(z)$. Anti-holomorphic functions satisfy $f_z(z) = 0$.

Injectivity and distortion. A continuously differentiable mapping f is locally injective and orientation-preserving if the following condition holds everywhere

$$|f_z| > |f_{\bar{z}}|, \quad (4)$$

where the inequality must be strict. It is easy to see from (4) that the Wirtinger derivative f_z cannot vanish, hence $f_z \neq 0$ is a necessary condition for local injectivity. It is also a sufficient condition if f is holomorphic. A common misconception is to identify complex holomorphic functions with planar conformal mappings. We emphasise that a holomorphic function induces a conformal mapping only if its derivative f' does not vanish, and that holomorphic functions in general (e.g. $f(z) = z^2$) are not locally injective.

It is often desired to measure and control the amount of angle and metric distortion induced by f . These can be formulated in terms of the singular values $0 \leq \sigma_f \leq \Sigma_f$ of the Jacobian matrix J_f , which can be expressed by

$$\Sigma_f = |f_z| + |f_{\bar{z}}|, \quad \sigma_f = \left| |f_z| - |f_{\bar{z}}| \right|. \quad (5)$$

If f is locally injective orientation-preserving, the latter expression simplifies to $\sigma_f = |f_z| - |f_{\bar{z}}|$. The quantity $\mu_f = \frac{f_{\bar{z}}}{f_z}$ is called the *first complex dilatation* and its modulus $k_f = \frac{|f_{\bar{z}}|}{|f_z|}$ is the *little dilatation*, which is related to the *large dilatation* K_f (the ratio of singular values) by $K_f = \frac{\Sigma_f}{\sigma_f} = \frac{1+k_f}{1-k_f}$. Both k_f and K_f measure the amount of conformal distortion and they are positive and monotonic functions of each other. If f is conformal then $k_f = 0$ and $K_f = 1$ at every point in the domain.

The Hilbert transform. The Hilbert transform [Bell 1992] is a linear operator that takes a harmonic function $u(z)$ and computes a second harmonic function $v(z)$ such that $u(z) + iv(z)$ is holomorphic. The pair u, v are called harmonic conjugates. The Hilbert transform always exists when the domain is simply-connected and is unique up to an additive imaginary constant.

3.1 Convexity

We conclude this section with some facts on convexity which will be useful in the next sections. An inequality constraint of the form $f(x_1, \dots, x_n) \leq 0$ defines a convex region in \mathbb{R}^n if f is a convex function. More generally, $f(x_1, \dots, x_n) \leq g(x_1, \dots, x_n)$ defines a convex region if f is a convex function and g is a concave function. Many useful functions are convex. For example, norms $\|x\|_p$ are convex and linear/affine functions are both convex and concave. Convexity and concavity are preserved under some operations. For example, we have that the positive weighted sum of convex (concave) functions is convex (concave). We also have that composition with an affine function preserves convexity and concavity. Meaning, if f is convex (concave) and h is affine, then $f(h)$ is also convex (concave). For further reading on convexity we refer to [Boyd and Vandenberghe 2004; Grant et al. 2008].

4 Convex Harmonic Characterization

Definition 1. A planar harmonic mapping $f : \Omega \subset \mathbb{C} \rightarrow \mathbb{C}$ is a (k, Σ, σ) bounded distortion mapping if it satisfies the following conditions

$$k_f(z) = \frac{|f_{\bar{z}}(z)|}{|f_z(z)|} \leq k < 1 \quad \forall z \in \Omega, \quad (6a)$$

$$\Sigma_f(z) = |f_z(z)| + |f_{\bar{z}}(z)| \leq \Sigma < \infty \quad \forall z \in \Omega, \quad (6b)$$

$$0 < \sigma \leq |f_z(z)| - |f_{\bar{z}}(z)| = \sigma_f(z) \quad \forall z \in \Omega, \quad (6c)$$

where k, Σ, σ are real constants.

Equation (6a) sets an upper bound on the dilatation (conformal distortion) while (6b), (6c) set upper and lower bounds on the singular values of J_f . Together, (6b), (6c) imply that common measures of isometric distortion such as (but not limited to) $\tau_f =$

$\max(\Sigma_f, 1/\sigma_f)$ are also bounded. We further note that a bounded distortion mapping is locally injective and orientation-preserving (see [Chen and Weber 2015], Observation 3).

Definition 2 (\mathcal{BD} space). Given a simply-connected domain $\Omega \subset \mathbb{C}$, the \mathcal{BD} space is the function space that contains all the (k, Σ, σ) bounded distortion harmonic mappings $f : \Omega \rightarrow \mathbb{C}$ (for some choice of the constants k, Σ, σ).

Searching for a particular mapping f within \mathcal{BD} by means of mathematical optimization is highly challenging since it is not a convex space. To illustrate the nonconvexity, consider the identity mapping $Id(z) = z$, and $rot(z) = -z$ which rotates vectors by π . Id and rot are holomorphic, hence harmonic, and they are both isometries, hence bounded distortion. Yet, their average $\frac{1}{2}Id + \frac{1}{2}rot$ is the zero mapping which is clearly not in \mathcal{BD} .

In this section, we show that there exists an alternative convex space (denoted \mathcal{H}) that is in one-to-one correspondence with \mathcal{BD} . This fact will be utilized in Section 5 to project an arbitrary reference mapping to \mathcal{BD} by solving a convex optimization problem. Intuitively, our alternative space \mathcal{H} relies on the fact that any harmonic mapping f (on a simply-connected domain) can be decomposed into a sum of holomorphic and anti-holomorphic functions $f = \Phi + \bar{\Psi}$. Such a decomposition always exists and it is unique up to a constant that can be fixed by setting $\Psi(z_0) = 0$ at an arbitrary anchor point $z_0 \in \Omega$ (see [Duren 2004] Section 1.2 for a proof). The other direction is also true. Namely, that a sum of any holomorphic and anti-holomorphic functions is harmonic. The anti-holomorphic part $\bar{\Psi}$ will be shared among the \mathcal{BD} and \mathcal{H} spaces, and the holomorphic part Φ will be parameterized in \mathcal{H} through the boundary values of the real-valued function $r = |\Phi'|$ [Weber and Gotsman 2010]. This is possible since we are only interested in locally injective harmonic mappings, for which the holomorphic part Φ is in fact a conformal mapping (Lemma 12 in Appendix A). Let us start by formally defining \mathcal{H} .

Definition 3 (\mathcal{H} space). Let $\Omega \subset \mathbb{C}$ be a simply-connected domain in the complex plane. Let $r(w) : \partial\Omega \rightarrow \mathbb{R}$ be a continuous real-valued function defined on the boundary of the domain and $\Psi(z) : \Omega \rightarrow \mathbb{C}$ be holomorphic in Ω .

A pair of functions $h = \{\Psi(z), r(w)\}$ belongs to the function space \mathcal{H} if the following inequalities are satisfied at every boundary point

$$|\Psi'(w)| \leq k r(w) \quad \forall w \in \partial\Omega, \quad (7a)$$

$$|\Psi'(w)| \leq \Sigma - r(w) \quad \forall w \in \partial\Omega, \quad (7b)$$

$$|\Psi'(w)| \leq r(w) - \sigma \quad \forall w \in \partial\Omega. \quad (7c)$$

The constants k, Σ, σ , bound the distortion and we assume that

$$0 < \sigma \leq \Sigma < \infty, \quad 0 \leq k < 1.$$

Ψ' is the first complex derivative of Ψ , and $|\cdot|$ is the modulus. Furthermore, we augment the definition by some constants: $z_0 \in \Omega$ is a user defined anchor point where we assume $\Psi(z_0) = 0$. $\theta \in (-\pi, \pi]$, and $d \in \mathbb{C}$ are additional constants.

The constants are merely needed to make the representation unique. $z_0 \in \Omega$ is a reference point where we can fix a global translation ($d \in \mathbb{C}$) and a rotation ($\theta \in (-\pi, \pi]$) which are degrees of freedom that arise from integration. As we will show next, \mathcal{H} characterizes the space of bounded distortion harmonic mappings and, somewhat surprisingly, unlike \mathcal{BD} it is convex as proven below.

Proposition 4 (Convexity). The \mathcal{H} space is convex.

Proof. Based on the arguments given in Section 3.1, the right hand sides of (7) are affine functionals, hence concave. Additionally, the left hand side of (7) is a convex functional since $|\cdot|$ is convex and the derivative is a linear operator. More formally, for any two holomorphic functions Ψ_0, Ψ_1 , and $\forall t \in [0, 1]$ we have

$$|((1-t)\Psi_0 + t\Psi_1)'| = |(1-t)\Psi_0' + t\Psi_1'| \leq (1-t)|\Psi_0'| + t|\Psi_1'|,$$

where the right inequality is due to the triangle inequality and the nonnegativity of t and $1-t$. \square

Optimizing in the convex space \mathcal{H} is easier computationally, however, efficient operators that transfer mappings from \mathcal{BD} to \mathcal{H} and vice versa are also needed. These operators are presented next.

Definition 5. $\mathcal{F} : \mathcal{BD} \rightarrow \mathcal{H}$ is an operator that takes a mapping $f \in \mathcal{BD}$ to \mathcal{H} using the following procedure. First, decompose f as follows

$$f(z) = \Phi(z) + \overline{\Psi}(z), \quad \Psi(z_0) = 0 \quad (8)$$

where Φ and Ψ are holomorphic functions and $z_0 \in \Omega$ is an arbitrary anchor point in the domain. Then, the corresponding point in \mathcal{H} is defined by $\{\Psi(z), r(w)\} = \{\Psi(z), |\Phi'(w)|\}$, i.e., Ψ is maintained, and the function r is defined by differentiating Φ , applying $|\cdot|$ and restricting to the boundary. Finally, the constants are determined by: $d = f(z_0), \theta = \text{Arg}(\Phi'(z_0))$, where Arg is the principal branch of the complex argument.

As can be seen, evaluating \mathcal{F} is straightforward. Let us obtain some intuition about the function r . By decomposing the harmonic mapping f into its holomorphic and anti-holomorphic parts (Equation (8)), we can obtain the following simple expressions for the Wirtinger derivatives (Section 3): $f_z = \overline{\Psi}'$ and $f_{\bar{z}} = \Phi'$. Taking norm on both sides gives $|f_z| = |\overline{\Psi}'|$ and $|f_{\bar{z}}| = |\Phi'| = r$ from which we see that the function r in the alternative \mathcal{H} space represents $|f_z(w)|$. Rearranging the nonconvex constraints of \mathcal{BD} (6) gives

$$|f_z(z)| \leq k |f_z(z)| \quad \forall z \in \Omega, \quad (9a)$$

$$|f_z(z)| \leq \Sigma - |f_z(z)| \quad \forall z \in \Omega, \quad (9b)$$

$$|f_z(z)| \leq |f_z(z)| - \sigma \quad \forall z \in \Omega. \quad (9c)$$

Let us verify that $\mathcal{F}(f)$ indeed belongs to \mathcal{H} . This is simply achieved by substituting $|f_z|$ with $|\overline{\Psi}'|$, and $|f_{\bar{z}}|$ with r in the equations above and restricting to the boundary which results in (7). The intuition for our construction of the \mathcal{H} space is that in the above nonconvex inequalities, the lhs is convex and the rhs of (9a),(9c) are also convex. Substituting $|f_z|$ with the affine expression r , ensures that the rhs is affine, hence concave (Section 3.1) and all three inequalities are convex.

Going from \mathcal{H} back to \mathcal{BD} is slightly more involved but is still straightforward and can be done with the help of the (linear) Hilbert transform.

Definition 6. $\mathcal{F}^{-1} : \mathcal{H} \rightarrow \mathcal{BD}$ is an operator that maps a pair of functions $h = \{\Psi(z), r(w)\} \in \mathcal{H}$ and constants $z_0 \in \Omega, \theta \in (-\pi, \pi], d \in \mathbb{C}$ to a harmonic mapping in \mathcal{BD} using the following procedure. First, we evaluate $\ln(r(w))$ on $\partial\Omega$ and extend it harmonically to the interior of Ω by solving the Dirichlet problem, obtaining $\xi(z)$. Using the Hilbert transform, we obtain a harmonic conjugate function $\zeta(z)$ which is unique for the choice $\zeta(z_0) = \theta$. The function $l(z) = \xi(z) + i\zeta(z)$ is holomorphic and so is $\Phi(z) = \int e^{l(z)}$ (where \int denotes the antiderivative). The antiderivative is uniquely defined by setting $\Phi(z_0) = d$. Finally, $\mathcal{F}^{-1}(h)$ is given by the sum $f = \Phi + \overline{\Psi}$.

Lemma 14 in the Appendix asserts that the operator \mathcal{F}^{-1} is well-defined and that it always produces a unique harmonic mapping f with bounded distortion. Moreover, the Lemma justifies the notation for \mathcal{F}^{-1} by showing that it is the inverse of the operator \mathcal{F} . We are now ready to state the main theorem which is formally proved in Appendix A.

Theorem 7. $\mathcal{F} : \mathcal{BD} \rightarrow \mathcal{H}$ is a one-to-one correspondence (bijection).

The practical meaning of Theorem 7 is that one can search for a candidate h in the convex space \mathcal{H} and that for any outcome, a unique bounded distortion harmonic mapping f can be easily recovered using \mathcal{F}^{-1} . Moreover, any mapping $f \in \mathcal{BD}$ has its representative $h \in \mathcal{H}$, so we do not “lose” any mapping by using the alternative \mathcal{H} space.

We have shown that by a suitable change of variables, the nonconvex \mathcal{BD} space can be transformed into a convex one. However, in order to have a convex optimization, the objective function needs to be convex as well. This precludes for example, the ability to directly control position or orientation which are nonlinear in the alternative variables of \mathcal{H} . Handling these is addressed in Section 6 through the definition of two additional alternative spaces. These spaces are not convex and are theoretically less appealing, nonetheless, their near-convexity and guaranteed feasibility leads to superior results.


In the next section we will discretize \mathcal{H} and explain how to project an arbitrary reference mapping g to the bounded distortion harmonic mappings space by solving a convex optimization problem.

5 Optimization

Given user prescribed bounds $0 \leq k < 1$, $0 < \sigma \leq \Sigma$, and an arbitrary reference mapping g , we design a convex constrained optimization for the projection of g onto \mathcal{BD} by minimizing the following convex energy functional

$$\oint_{\partial\Omega} (r(w) - |g_z(w)|)^2 ds + \lambda_{\mathcal{H}} \iint_{\Omega} |\Psi(z)' - \overline{g_z}(z)|^2 da. \quad (10)$$

Assuming that the distortion of g exceeds the user specified bounds (i.e., $g \notin \mathcal{BD}$), our algorithm first transfers g to the alternative representation using the same \mathcal{F} of Definition 5. $|g_z|$ in the energy above can be thought of as r_g , and $\overline{g_z}$ plays the role of $\overline{\Psi}'_g$. The optimization projects $\{\Psi_g, r_g\}$ to a point $\{\Psi, r\}$ in \mathcal{H} which is then transformed by \mathcal{F}^{-1} to a point in \mathcal{BD} corresponding to a bounded distortion mapping f . The inset provides a schematic illustration.



The unique property of our algorithm is that regardless of how badly distorted the reference mapping g is, and for any choice of the bounds (strict as they can be), the algorithm always converges and produces a valid mapping that brings the (discrete) energy to its unique global minimum.

5.1 Discretization

Let P be a user defined simply-connected polygon oriented counterclockwise. Our domain Ω is defined as the interior of P . Let $\{z_1, z_2, \dots, z_n\}$ be the vertices of a polygon \hat{P} (the cage) that is obtained by slightly offsetting P in the outward normal direction, and let $z_0 \in \Omega$ be a user defined anchor point. We sample the boundary P uniformly to obtain the set \mathcal{A} . The C^∞ mappings we compute are visualized on a mesh \mathcal{M} that is obtained by triangulating P . The set

\mathcal{B} consists of the triangle barycenters of \mathcal{M} . We used a mesh with roughly 10,000 triangles for all the results presented in the paper.

The holomorphic function Ψ is defined based on the discrete Cauchy transform [Weber et al. 2009]

$$\Psi(z) = \sum_{j=1}^n C_j(z) \psi_j, \quad (11)$$

where $C_j(z) \in \mathbb{C}$ is the j^{th} holomorphic Cauchy complex barycentric coordinate associated with vertex z_j , and ψ_j are complex coefficients. $C_j(z)$ and its complex derivative, $C'_j(z)$, possess rather simple closed-form expressions [Weber 2010, Appendices B,C]. For any choice of the coefficients ψ_j , the conjugate of the derivative of the anti-holomorphic part is

$$\overline{f'_z}(z) = \Psi'(z) = \sum_{j=1}^n C'_j(z) \psi_j, \quad (12)$$

and is precisely holomorphic. Due to the offset being used, the derivatives in (12) can be easily evaluated at any point z inside or on P using a simple formula. We typically use an offset of 0.1% of the overall length of P . Finally, the function $r(w)$ is discretized as a continuous piecewise affine function on the boundary and is fully determined by r_i , its value at the samples in \mathcal{A} . In order to enrich the finite-dimensional holomorphic subspace of Equation (11), we also super sample \hat{P} such that it has at least 25 vertices.

5.2 The Projection Algorithm

We first evaluate $|g_z(p_i)|$ and $\overline{g_z}(p_i)$ on the samples of \mathcal{A} and \mathcal{B} respectively. If g is a harmonic mapping expressed in the Cauchy basis, then $g_z, g_{\bar{z}}$ have closed-form expressions. Otherwise, if g is a piecewise affine mapping defined on \mathcal{M} , we compute $g_z, g_{\bar{z}}$ using the formula provided in [Weber et al. 2012, Equation (6)]. If g has a different triangulation than \mathcal{M} , we simply use the triangle that contains the sample p_i .

In the second step, we project $|g_z|$ and $\overline{g_z}$ to \mathcal{H} by solving the following convex optimization problem

$$\begin{aligned} \min_{\psi_1 \dots \psi_n, r_1 \dots r_{|\mathcal{A}|}} \quad & E_{\mathcal{H}} \\ \text{s.t.} \quad & \Psi(z_0) = 0, \\ & \forall p_i \in \mathcal{A} \quad |\Psi'(p_i)| \leq k r_i, \\ & \forall p_i \in \mathcal{A} \quad |\Psi'(p_i)| \leq \Sigma - r_i, \\ & \forall p_i \in \mathcal{A} \quad |\Psi'(p_i)| \leq r_i - \sigma. \end{aligned} \quad (13)$$

$E_{\mathcal{H}}$ is obtained by approximating the integral in Equation (10) with the following quadratic function

$$E_{\mathcal{H}} = \sum_{i=1}^{|\mathcal{A}|} (r_i - |g_z(p_i)|)^2 + \lambda_{\mathcal{H}} \sum_{i=1}^{|\mathcal{B}|} |\Psi'(p_i) - \overline{g_z}(p_i)|^2, \quad (14)$$

where $\lambda_{\mathcal{H}}$ is a positive weight that balances the two energy terms. We used $\lambda_{\mathcal{H}} = 1$ for all our experiments.

Now that we have a point $h \in \mathcal{H}$, the third step of our algorithm approximates the \mathcal{F}^{-1} operator (Definition 6). We evaluate $\ln(r_i)$ on the boundary samples of \mathcal{A} , then compute the harmonic extension to $\ln(r)$ (Dirichlet problem) and its harmonic conjugate (Hilbert transform) simultaneously based on the boundary element method (BEM) described in [Weber and Gotsman 2010] Section 6. The idea is to construct a holomorphic function $l(z) = \sum_{j=1}^n C_j(z) l_j$




Figure 3: Projection of a holomorphic function generated by the Cauchy coordinates [Weber et al. 2009] via the \mathcal{H} space with an upper bound $k = 0$, and a lower bound $\sigma = 0.2$, such that the mapping becomes conformal.

and search for the complex coefficients l_j such that $\text{Re}(l(z))$ on the boundary is as close as possible to the prescribed boundary values. This boils down to solving the following least-squares problem

$$\begin{aligned} \min_{l_1 \dots l_n} \quad & \sum_{i=1}^{|\mathcal{A}|} \left(\text{Re}(l(p_i)) - \ln(r_i) \right)^2 \\ \text{s.t.} \quad & \text{Im}(l(z_0)) = \text{Arg}(g_z(z_0)). \end{aligned} \quad (15)$$

The equality constraint nails down the constant degree of freedom of the harmonic conjugate function. Geometrically, it corresponds to a global rotation of f_z and is fixed based on the corresponding value at g (Definitions 5 and 6).

The function $l(z)$ represents the logarithm of $f_z(z)$, hence once the variables l_j are found, we can evaluate $\Phi'(z) = e^{l(z)}$ at any point $z \in \Omega$. In order to complete the process and obtain $f = \Phi + \Psi$, we need the antiderivative of Φ' . The closed-form expression we obtained for Φ' is holomorphic, hence it is *precisely* integrable. However, we could not obtain a closed-form expression for its antiderivative. Nevertheless, we have found that since the Cauchy coordinates are smooth and Lipschitz (see [Chen and Weber 2015] for their Lipschitz constants), numerical integration provides superior results with almost no computational overhead (100ms on our most complicated model).

Algorithm 1 Projection via the \mathcal{H} space

- 1: **Input:** $g \notin \mathcal{BD}, z_0 \in \Omega, k, \Sigma, \sigma$ **Output:** $f \in \mathcal{BD}$
 - 2: evaluate $|g_z|, \overline{g_z}$
 - 3: $\{\Psi(z), r\} := \text{solve (13)}$
 - 4: $l(z) := \text{solve (15)}$
 - 5: $\Phi'(z) := e^{l(z)}$
 - 6: $\Phi(z) := \int \Phi'$ such that $\Phi(z_0) = g(z_0)$
 - 7: $f := \Phi(z) + \Psi(z)$
-

The integration is done using the trapezoidal rule along a spanning tree of the graph induced by \mathcal{M} , starting from the anchor point z_0 as its root, for which we set $\Phi(z_0) = g(z_0)$ for uniqueness. This choice ensures that the input reference mapping agrees with the projected mapping at z_0 , i.e. $f(z_0) = g(z_0)$. Any other choice would correspond to a global translation of f . Algorithm 1 provides a summary of the projection to \mathcal{BD} via the convex \mathcal{H} space.

Figures 3,4,5 show the result of our projection algorithm and compare it with the methods of [Lipman 2012; Aigerman and Lipman 2013; Kovalsky et al. 2015; Chen and Weber 2015].

We showed how to efficiently compute a bounded distortion mapping that is similar to a reference mapping g , where similarity is measured by Equation (10). While this procedure is optimal in the sense that it obtains the closest $f \in \mathcal{BD}$ to g , the similarity measure that we use for Φ , the conformal part of f , is based solely on $|g_z|$ and is oblivious to the argument of g_z . Intuitively, g_z represents the closest similarity transformation to the Jacobian J_f and its argument represents the closest rotation. In the next section, we

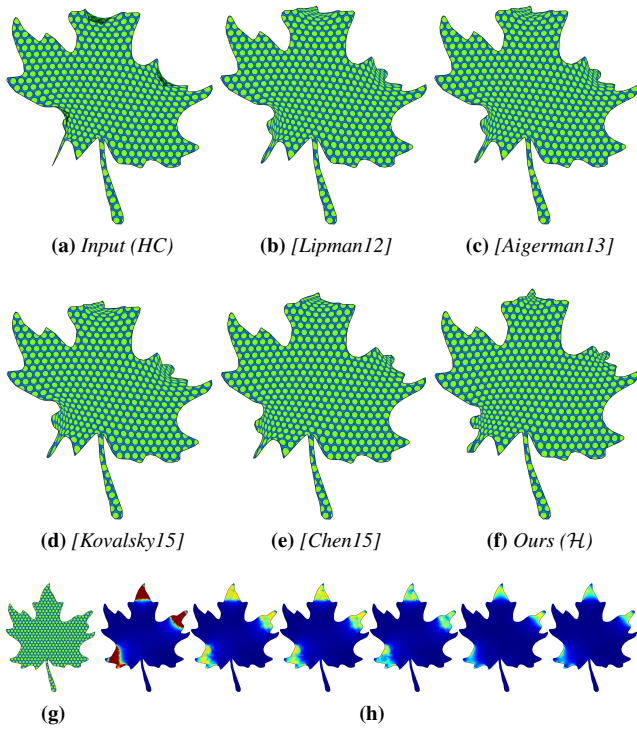


Figure 4: Leaf. (g) Domain. (a) Highly distorted reference mapping with large areas in which the orientation is reversed, generated using harmonic coordinates [Joshi et al. 2007]. (h) Visualization of the isometric distortion $\tau_f = \max(\Sigma_f, 1/\sigma_f)$.

introduce two additional spaces that characterize Φ in a different manner. Rather than using the real-valued function $|\Phi'(w)|$ on $\partial\Omega$, they are based on the holomorphic function $\log \Phi' : \Omega \rightarrow \mathbb{C}$

$$\log \Phi'(z) = \ln |\Phi'(z)| + i \arg(\Phi'(z)).$$

Such a representation captures both the scale and the rotational part of Φ . In contrast to \mathcal{H} , these new spaces are not convex. Nevertheless, we can still optimize in these spaces using convex optimization and with the same mathematical guarantees for *convergence* and *feasibility* that \mathcal{H} provides. The mappings obtained using these new projection operators tend to better resemble (visually) the reference mappings and are arguably better qualitatively. The first logarithmic space, \mathcal{L}_ν (Section 6.1) is shown to be nearly convex, while the second one, \mathcal{L}_ψ (Section 6.2) supports positional constraints.

6 Convex Logarithmic Characterization

Any harmonic mapping (not necessarily bounded distortion) can be parameterized by its holomorphic derivatives $f_z, \overline{f_z}$, since for any such functions, their antiderivatives Φ, Ψ exist and are also holomorphic, such that the sum $f = \Phi + \Psi$ is harmonic. We already saw (proved formally in Lemma 12 - Appendix A) that a necessary condition for local injectivity of a harmonic mapping f is that Φ is conformal (have nonvanishing f_z). This requirement can be expressed by $|f_z| > 0$ (the inequality is strict), but since this is a nonconvex constraint, it is hard to enforce in practice. Instead, it will be beneficial to parameterize f_z using its logarithm [Weber and Gotsman 2010]. With such a parameterization, no constraints are required. To see this, let $l(z)$ be any holomorphic function. Then, $e^{l(z)}$ is also holomorphic and nonvanishing, and it can be used to represent the derivative f_z of a conformal mapping.

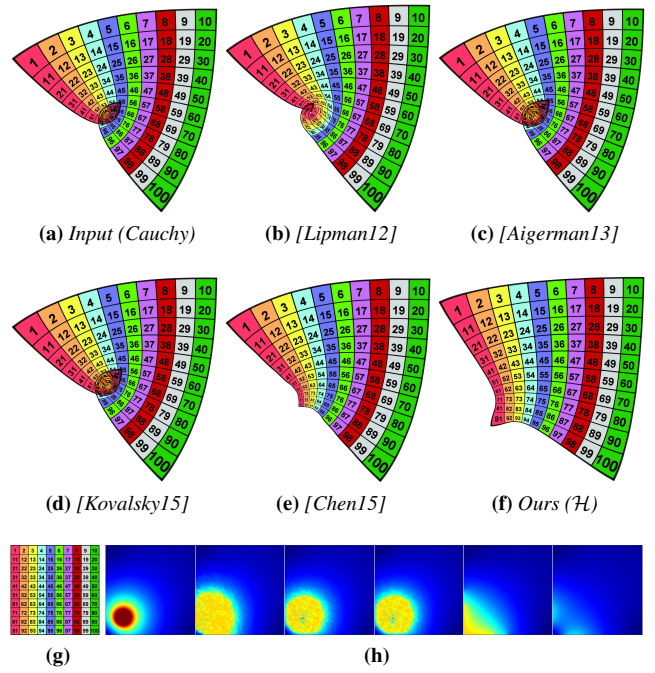


Figure 5: Square Loop. (g) Domain. Note the symmetry in our result. The isometric distortion (visualized) show that the input mapping has vanishing g_z inside the domain (red area).

6.1 The Logarithmic \mathcal{L}_ν Space

The *second complex dilatation* of a complex-valued function f is defined as $\nu = \frac{\overline{f_z}}{f_z}$ and it is related to the first complex dilatation μ by $k = |\nu| = |\mu|$. For the special case of locally injective harmonic complex-valued functions, ν is holomorphic. Additionally, since $\overline{f_z} = f_z \nu$, one can fully characterize harmonic mappings based on f_z and ν rather than with f_z and $\overline{f_z}$.

The \mathcal{L}_ν space is based on two holomorphic functions $l(z) = \log f_z$ and $\nu(z) = \frac{\overline{f_z}}{f_z} = \frac{\overline{f_z}}{e^l}$. Let us express Σ_f in terms of l and ν

$$\Sigma_f = |f_z| + |f_z| = |f_z| (1 + |\nu|) = e^{\text{Re}(l)} (1 + |\nu|).$$

Similarly $\sigma_f = e^{\text{Re}(l)} (1 - |\nu|)$, and with that we are ready to introduce the \mathcal{L}_ν space.

Definition 8 (\mathcal{L}_ν space). Let $l(z), \nu(z) : \Omega \rightarrow \mathbb{C}$ be two holomorphic functions defined on a simply-connected domain Ω . A pair $\{l(z), \nu(z)\}$ belongs to the function space \mathcal{L}_ν if the following inequalities are satisfied at every boundary point

$$k_f(w) = |\nu(w)| \leq k \quad \forall w \in \partial\Omega, \quad (16a)$$

$$\Sigma_f(w) = e^{\text{Re}(l(w))} (1 + |\nu(w)|) \leq \Sigma \quad \forall w \in \partial\Omega, \quad (16b)$$

$$\sigma \leq e^{\text{Re}(l(w))} (1 - |\nu(w)|) = \sigma_f(w) \quad \forall w \in \partial\Omega. \quad (16c)$$

Just like the \mathcal{H} space, \mathcal{L}_ν has one-to-one correspondence with \mathcal{BD} (under suitable choice of some integration constants). The mechanism of the proof is quite similar to that of Theorem 7 and is omitted here for brevity. Given a mapping $f \in \mathcal{BD}$, the operator $\mathcal{F} : \mathcal{BD} \rightarrow \mathcal{L}_\nu$ is defined by $\{l(z), \nu(z)\} = \{\log f_z(z), \frac{\overline{f_z(z)}}{f_z(z)}\}$, where $\log f_z$ is a holomorphic branch which exists since f_z does not vanish. For such a branch to be unique, we set $\text{Im}(\log f_z(z_0)) = \text{Arg}(f_z(z_0))$ for some arbitrary anchor point z_0 and augment the definition with a constant $d \in \mathbb{C}$ such that $d = f(z_0)$.

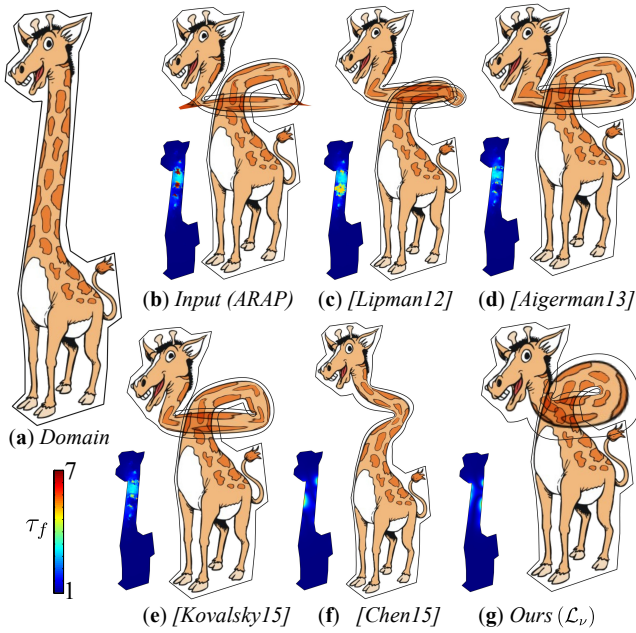


Figure 6: Giraffe. (b,d,e) are not locally injective. (c) is bounded distortion but nonsmooth. (c,f) have wrong orientation for the neck. (g) Smooth, bounded distortion and visually resembles the input.

The inverse operator $\mathcal{F}^{-1} : \mathcal{L}_\nu \rightarrow \mathcal{BD}$ is given by $f = \Phi + \bar{\Psi}$, where $\Phi = \int e^{l(z)}$ and $\Psi = \int \nu(z)e^{l(z)}$. The degrees of freedom in the antiderivatives are chosen such that $\Phi(z_0) = d$ and $\Psi(z_0) = 0$. This ensures $f(z_0) = d$ and conforms with the definition of \mathcal{F} .

Condition (16a) bounds the conformal distortion and is convex since the norm is convex. Condition (16c) bounds σ_f from below and can be rearranged as $\sigma e^{-\text{Re}(l)} + |\nu| \leq 1$, which is also convex since the exponent is convex, hence the left-hand side is convex (sum of convex). Condition (16b) bounds Σ_f from above and unfortunately, is not convex. Nonetheless, we can approximate it using a convex condition without losing much, in a way that bounds the original function and without compromising feasibility. Let us divide both sides of (16b) by $e^{\text{Re}(l(w))}$ and take \ln on both sides

$$\ln(1 + |\nu(w)|) \leq \ln \Sigma - \text{Re}(l(w)). \quad (17)$$

Next, decompose (17) into two conditions by introducing an auxiliary function $s(w)$


$$s(w) = \ln \Sigma - \text{Re}(l(w)) \quad \forall w \in \partial\Omega, \quad (18a)$$

$$\ln(1 + |\nu(w)|) \leq s(w) \quad \forall w \in \partial\Omega, \quad (18b)$$

where (18a) is convex (affine) and (18b) is nonconvex. Condition (18b) is now the only nonconvex condition we have.

At a single point w , the boundary of the domain that is defined by (18b) can be visualized in \mathbb{R}^3 as a surface of revolution (i.e. with circular cross sections) with a concave generating curve (see the red curve in the inset). From (16c) we have that s is bounded from above by the plane $s = \ln(\Sigma/\sigma)$ and from below by the cylinder $|\nu| \leq k \leq 1$. As the inset shows, the concavity of this domain is very subtle. We substitute this domain with a second order convex cone (cyan surface)

$$|\nu(w)| \leq s(w). \quad (19)$$



Another way to derive the convexification is to consider the first term in the Taylor series expansion of $\ln(1 + |\nu|)$ around zero. The difference between the convex and nonconvex domains is marked with purple. The image illustrates the worse case scenario where the bound k chosen by the user is 1. Typically, k is much smaller and so, the difference between the domains becomes quite insignificant. By combining (18a) with (19) and adding the other two convex conditions, we arrive at the *convex* \mathcal{L}_ν space definition

$$|\nu(w)| \leq k \quad \forall w \in \partial\Omega, \quad (20a)$$

$$|\nu(w)| \leq \ln \Sigma - \text{Re}(l(w)) \quad \forall w \in \partial\Omega, \quad (20b)$$

$$\sigma e^{-\text{Re}(l(w))} + |\nu(w)| \leq 1 \quad \forall w \in \partial\Omega. \quad (20c)$$

The advantage of (20) over the convexification approach of [Lipman 2012] is threefold. First, it is *parameterless* and does not depend on the definition of local frames. Second, [Lipman 2012] convexifies the constraints for both k_f and σ_f while we only convexify the constraint for Σ_f . This is important, since our experiments show that the constraint for Σ_f rarely becomes active (implying that the solution we obtain is mostly optimal). Moreover, in applications where only local injectivity is required, it can be dispensed altogether. Third, the above space is nonempty for any user specified bounds (as strict as they can be), which means that the corresponding optimization problem is always *feasible*. This is formally expressed below (see Appendix B for a proof).

Proposition 9. For any real constants k, Σ, σ satisfying

$$0 < \sigma \leq \Sigma < \infty, \quad 0 \leq k < 1,$$

there exist holomorphic $l(z), \nu(z) : \Omega \rightarrow \mathbb{C}$ satisfying (20).

Implementation details. Similarly to Section 5.1, we discretize the holomorphic functions $l(z)$ and $\nu(z)$ with the Cauchy transform

$$l(z) = \sum_{j=1}^n C_j(z)l_j, \quad \nu(z) = \sum_{j=1}^n C_j(z)\nu_j. \quad (21)$$

The projection of the reference mapping g on \mathcal{L}_ν is done by solving the following convex optimization problem

$$\begin{aligned} \min_{l_1 \dots l_n, \nu_1 \dots \nu_n} \quad & E_l + \lambda_\nu E_\nu \\ \text{s.t.} \quad & \forall p_i \in \mathcal{A} \quad |\nu(p_i)| \leq k, \\ & \forall p_i \in \mathcal{A} \quad |\nu(p_i)| \leq \ln \Sigma - \text{Re}(l(p_i)), \\ & \forall p_i \in \mathcal{A} \quad \sigma e^{-\text{Re}(l(p_i))} + |\nu(p_i)| \leq 1, \end{aligned} \quad (22)$$

where

$$E_l = \sum_{i=1}^{|\mathcal{A}|} \left| \text{Re}(l(p_i)) - \ln(|g_z(p_i)|) \right| + \left| \text{Im}(l(p_i)) - \theta_i \right|, \quad (23)$$

and

$$E_\nu = \sum_{i=1}^{|\mathcal{B}|} \left| \nu(p_i) - \nu_g(p_i) \right|. \quad (24)$$

λ_ν is a weight that balances the two energy terms. We used $\lambda_\nu = 1$ for all the results. The energy E_l encourages $l(z)$ (the logarithm of f_z) to be as close as possible to the logarithm of g_z . θ_i are real constants representing the argument of g_z (we explain in Section 6.3 how to obtain them).

The energy E_ν encourages the second complex dilatation ν of the mapping f to be as close as possible to the one of g . Once the optimal solution is found we have a closed-form expression for

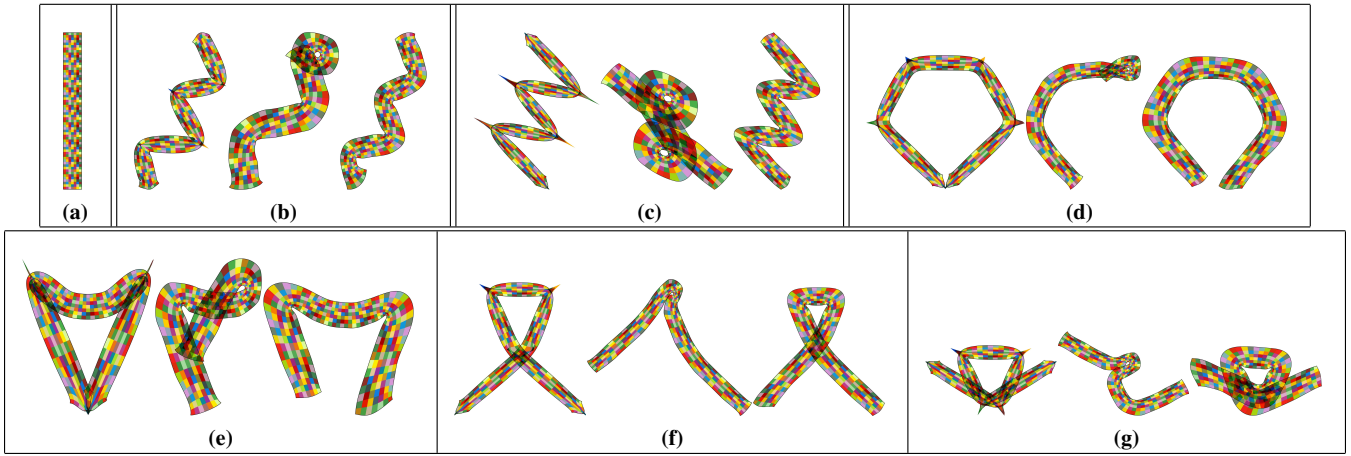


Figure 7: Colored bar. (a) Domain; (b-g) In each group from left to right: input mapping (ARAP), [Chen and Weber 2015], and ours (\mathcal{L}_ν).

$f_z(z) = e^{l(z)}$ and $\overline{f_{\bar{z}}}(z) = \nu(z) e^{l(z)}$ from which we compute $\Phi(z)$ and $\Psi(z)$ by numerical integration similarly to Section 5.2, where for uniqueness we set $\Phi(z_0) = g(z_0)$ and $\Psi(z_0) = 0$ at the user specified anchor points z_0 .

Algorithm 2 summarizes the steps of the projection based on the convex \mathcal{L}_ν space. Figures 6 and 7 show some results.

Algorithm 2 Projection via the \mathcal{L}_ν space

-
- 1: **Input:** $g \notin \mathcal{BD}$, $z_0 \in \Omega$, k, Σ, σ **Output:** $f \in \mathcal{BD}$
2: evaluate $g_z, \overline{g_{\bar{z}}}$ on \mathcal{A} and \mathcal{B}
3: evaluate $\ln |g_z|, \nu_g := \frac{\overline{g_{\bar{z}}}}{g_z}$
4: $\theta_g :=$ extract argument from g_z (Section 6.3)
5: $l(z), \nu(z) :=$ solve (22)
6: $\Phi(z) := \int e^{l(z)}$ such that $\Phi(z_0) = g(z_0)$
7: $\Psi(z) := \int \nu(z) e^{l(z)}$ such that $\Psi(z_0) = 0$
8: $f(z) := \Phi(z) + \overline{\Psi}(z)$
-

6.2 The Logarithmic \mathcal{L}_ψ Space

In some scenarios, it is also required to incorporate positional constraints into the optimization. Such constraints would be nonlinear in the variables of the \mathcal{L}_ν and \mathcal{H} spaces, and including them will render the optimization nonconvex. In general, projecting a mapping g to the bounded distortion space with more than one positional constraint $f(p_i) = q_i$ is not always possible since even the full nonconvex space would be empty for some choices of the constraints. The next space we introduce is designed to partially address this requirement by incorporating positional constraints in a soft fashion and without sacrificing feasibility.

Similarly to the \mathcal{L}_ν space, we use the function $l(z)$ to represent the logarithm of $f_z(z)$, but instead of $\nu(z)$ we use $\Psi(z)$ which, similarly to the \mathcal{H} space, represents the anti-holomorphic part of the harmonic mapping f . The bounds on the distortion can be expressed as usual in terms of $|f_z| = e^{\text{Re}(l)}$ and $|f_{\bar{z}}| = |\Psi'|$.

Definition 10 (\mathcal{L}_ψ space). Let $l(z), \Psi(z) : \Omega \rightarrow \mathbb{C}$ be two holomorphic functions defined on a simply-connected domain Ω . A pair $\{l(z), \Psi(z)\}$ belongs to the function space \mathcal{L}_ψ if the following inequalities are satisfied at every boundary point

$$|\Psi'(w)| \leq k e^{\text{Re}(l(w))} \quad \forall w \in \partial\Omega, \quad (25a)$$

$$e^{\text{Re}(l(w))} + |\Psi'(w)| \leq \Sigma \quad \forall w \in \partial\Omega, \quad (25b)$$

$$\sigma \leq e^{\text{Re}(l(w))} - |\Psi'(w)| \quad \forall w \in \partial\Omega. \quad (25c)$$

\mathcal{L}_ψ also fully characterizes \mathcal{BD} , albeit it is not a convex space (Conditions (25a), (25c) are not convex). Nonetheless, for any choice of fixed Ψ , these inequality constraints can be easily transformed to a convex form. However, how should we fix Ψ without the risk of getting an empty space? It turns out that by simply choosing $\Psi(z) = 0$ we can always find $l(z)$ that satisfies (25). Note that if $\Psi(z) = 0$, Equation (25a) is immediately satisfied and (25b) and (25c) take a simpler form

$$e^{\text{Re}(l(w))} \leq \Sigma \quad \forall w \in \partial\Omega, \quad (26a)$$

$$\sigma \leq e^{\text{Re}(l(w))} \quad \forall w \in \partial\Omega, \quad (26b)$$

Furthermore, (26) can be converted to the following equivalent convex (in fact linear) constraints by taking $\ln(\cdot)$ of both sides

$$\text{Re}(l(w)) \leq \ln(\Sigma) \quad \forall w \in \partial\Omega, \quad (27a)$$

$$\ln(\sigma) \leq \text{Re}(l(w)) \quad \forall w \in \partial\Omega. \quad (27b)$$

The convex space of (27) is nonempty due to the same consideration we applied in the proof of Proposition 9. Since $\Psi = 0$ implies $f_{\bar{z}} = 0$, these constraints actually fully characterize the space of conformal mappings on Ω with bounded isometric distortion. This is a convex subspace in the nonconvex \mathcal{L}_ψ space.

In order to project a reference mapping g to \mathcal{BD} via \mathcal{L}_ψ , we first compute a conformal mapping satisfying the constraints (27). Once $l(z)$ is found, we fix the function $e^{\text{Re}(l(w))}$ in (25) and let $\Psi(z)$ in (25) change freely. For fixed $l(z)$, the constraints in (25) become second order convex cones. Moreover, with the nonlinear part $l(z)$ fixed, positional constraints become linear and can be enforced.

Implementation details. The holomorphic function $l(z)$ is discretized, as usual as $l(z) = \sum_{j=1}^n C_j(z) l_j$. Our goal is to find a conformal mapping f with bounded isometric distortion such that the logarithm of f_z , which is represented by $l(z)$, is as close as possible to the logarithm of g_z , where g is the reference mapping. To this end, we solve the following convex optimization problem

$$\begin{aligned} \min_{l_1 \dots l_n} E_l \\ \text{s.t. } \forall p_i \in \mathcal{A} \quad \text{Re}(l(p_i)) \leq \ln(\Sigma), \\ \forall p_i \in \mathcal{A} \quad \ln(\sigma) \leq \text{Re}(l(p_i)), \end{aligned} \quad (28)$$

where E_l is the same energy used in (23). Once $l(z)$ is obtained, we fix $e^{\text{Re}(l(w))}$ in (25) and discretize it such that $r_i = e^{\text{Re}(l(p_i))}$

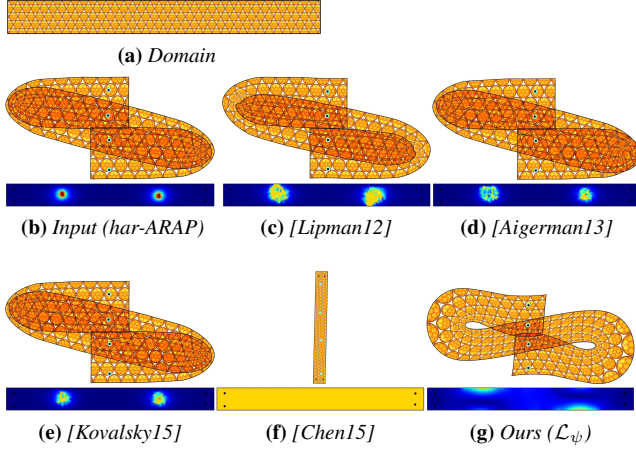


Figure 8: Orange bar with soft positional constraints. (d,e) resembles the reference (b) but are not locally injective. (c,f) are locally injective and obey the bounds but (c) is nonsmooth and (f) fails to satisfy the positional constraints.

are positive constants. $\Psi(z)$ is discretized as $\sum_{j=1}^n C_j(z)\psi_j$, and we solve the following convex optimization for the variables ψ_j

$$\begin{aligned} \min_{\psi_1 \dots \psi_n} \quad & \sum_{i=1}^{|B|} \left| \Psi'(p_i) - \overline{g_z}(p_i) \right|^2 + \lambda_{p2p} \sum_{i=2}^{|\mathcal{P}|} \left| \overline{f}(c_i) - \overline{f}_i \right|^2 \\ \text{s.t.} \quad & \Psi(c_1) = 0, \\ & \forall p_i \in \mathcal{A} \quad \left| \Psi'(p_i) \right| \leq k r_i, \\ & \forall p_i \in \mathcal{A} \quad \left| \Psi'(p_i) \right| \leq \Sigma - r_i, \\ & \forall p_i \in \mathcal{A} \quad \left| \Psi'(p_i) \right| \leq r_i - \sigma. \end{aligned} \quad (29)$$

The set $\mathcal{P} = \{c_i\}_{i=1}^{|\mathcal{P}|}$ contains points in Ω that the user wishes to map to corresponding target points f_i . These are incorporated into the energy as soft constraints (to ensure the problem is feasible). We used $\lambda_{p2p} = 1000$ for all the results. The inequalities are second order cone constraints and our energy is quadratic, since the term $\overline{f}(c_i)$ is linear in the variables ψ_j . The exact expression for it is obtained as follows. We numerically integrate (Section 5.2) the holomorphic function $e^{l(z)}$ such that $\Phi(c_1) = f_1$ and obtain $\Phi(z)$. This choice ensures that together with the constraint $\Psi(c_1) = 0$ in (29), the projected mapping f maps c_1 to f_1 . Next, note that $\Psi(c_i) = \sum_{j=1}^n C_j(c_i)\psi_j$ is linear in ψ_j , hence, $\overline{f}(c_i) = \overline{\Phi}(c_i) + \Psi(c_i)$ is linear.

Algorithm 3 summarizes the steps of the projection process.

Algorithm 3 Projection via the \mathcal{L}_ψ space

- 1: **Input:** $g \notin \mathcal{BD}$, k, Σ, σ , $c_i \rightarrow f_i$ **Output:** $f \in \mathcal{BD}$
2: evaluate $g_z, \overline{g_z}$ on \mathcal{A} and \mathcal{B}
3: evaluate $\ln |g_z|$
4: $\theta_g :=$ extract argument from g_z (Section 6.3)
5: $l(z) :=$ solve (28)
6: $\Phi(z) := \int e^{l(z)}$ such that $\Phi(c_1) = f_1$
7: $r_i := e^{\text{Re}(l(p_i))}$
8: $\Psi(z) :=$ solve (29)
9: $f(z) := \Phi(z) + \overline{\Psi}(z)$
-

6.3 Argument Extraction

The projection of a reference mapping g to the \mathcal{L}_ν and \mathcal{L}_ψ spaces requires as input the argument function θ_g of the derivative g_z ,

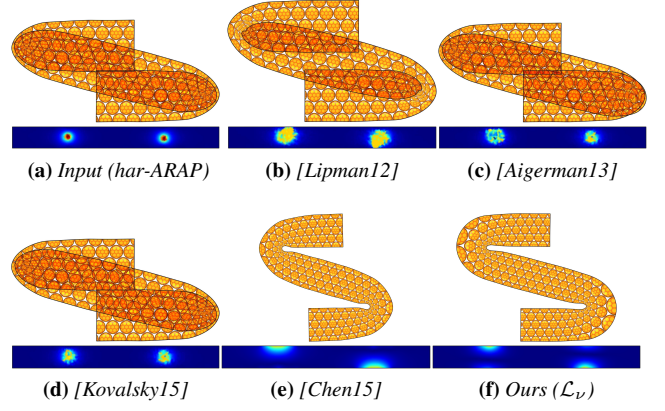


Figure 9: Orange bar without soft positional constraints.

which encodes the closest local rotation. Such an argument function $\theta_g(z)$ is a solution to the equation $\frac{g_z(z)}{|g_z(z)|} = e^{i\theta_g(z)}$. If g_z vanishes there is no solution. Otherwise, there are a countable number of solutions: $\theta_g(z) = \{\text{Arg}(g_z(z)) + 2\pi m(z) \mid m(z) \in \mathbb{Z}\}$, where Arg is a single-valued function that returns values in the range $(-\pi, \pi]$. Since the argument function is multi-valued, one has to choose $m(z)$ in order to obtain a single-valued function to feed the algorithm. One possible choice is $m(z) \equiv 0$, i.e., to use the principal branch Arg function. However, such a single-valued function is not necessarily continuous inside the domain.

Let us emphasize that our projection is guaranteed to produce a C^∞ mapping f that satisfies all the constraints, regardless of how badly distorted g is, including cases for which g or θ_g are not even continuous. However, a mapping f in \mathcal{BD} always has a continuous (in fact harmonic) branch of the argument θ_f . Hence, using θ_g , which is noncontinuous inside the domain, as a reference makes little sense.

An important question arises: Is it possible to find a single-valued argument function that (in contrast to Arg) is continuous in Ω ? Unfortunately, the answer is negative in cases where g_z vanishes inside the domain. For such cases, we construct a continuous (harmonic) function $\theta_g(z)$ that agrees with the argument of g_z on the boundary. The main idea is to push the discontinuity to the boundary of the domain at strategic points. Then, the continuity in the interior will be implied implicitly due to the harmonicity of $\theta_g(z)$.

We rely on the fact that for nonvanishing holomorphic functions, a continuous branch of the logarithm always exists, and it is holomorphic. The idea is to start by computing a holomorphic function $\Gamma(z)$ that is as close as possible to g_z (unless the input g is harmonic, and g_z is holomorphic already). Then, we search for the (isolated) zeros of $\Gamma(z)$ and in case they exist, we “push” them to the closest boundary point.

Projecting g_z to the holomorphic function space is done by solving a simple unconstrained least-squares problem

$$\min_{\gamma_1 \dots \gamma_n} \sum_{i=1}^{|\mathcal{A}|} \left| \Gamma(p_i) - g_z(p_i) \right|^2, \quad (30)$$

where $\Gamma(z)$ is discretized as $\sum_{j=1}^n C_j(z)\gamma_j$.

Next, we search Ω for zeros of $\Gamma(z)$ by iterating all the triangles in our triangulation, searching for a zero inside each triangle. To this end, we apply Cauchy’s argument principle which asserts that the number of zeros of a holomorphic function inside a simply-

connected domain is given by the following boundary integral

$$N = \frac{1}{2\pi i} \oint_{\partial T_j} \frac{\Gamma'(w)}{\Gamma(w)} dw, \quad (31)$$

where ∂T_j here represents the three edges of the j^{th} triangle. Evaluation of the integral is done numerically and can be avoided for the vast majority of the triangles by employing the sufficient conditions for the lack of zeros that are provided by [Chen and Weber 2015, Section 6.4], such that only if these fail, the integration is employed.

Equation (31) has an interesting geometric meaning. The integrand is the derivative of the function $\log(\Gamma)$. Starting from a boundary point w , integrating along the boundary in a counterclockwise direction measures the total change in the argument of Γ . Each zero in the domain adds an extra loop to the image of the boundary curve under the mapping, increasing its turning number by 1. For example, if there is one zero present, the total change in argument would be precisely 2π and the turning number of the *target* boundary curve would be 2 rather than 1.

Our strategy to counteract these extra unnecessary turns of the boundary is quite simple, yet proved to be very effective. In case a zero is found at z_0 , we find a point $b_0 \in \partial\Omega$ on the boundary that is closest to z_0 . We then compute θ_g by integrating the angular change $d\theta$ along the boundary, but when we reach b_0 , we subtract 2π from θ_g and continue to integrate till we get back to the starting point. Doing this for all zeros ensures that the total change in argument will be 0. Note that the function θ_g that we have constructed is discontinuous on the boundary but smoothness of the final mapping is guaranteed due to the smoothness of the Cauchy basis functions. Algorithm 4 summarizes the process of extracting θ_g .

Algorithm 4 Argument extraction

```

1: Input:  $g_z(z)$    Output:  $\theta_g$ 
2: if  $g$  is harmonic then
3:    $\Gamma(z) := g_z(z)$ 
4: else
5:    $\Gamma(z) := \text{solve (30) to project } g_z(z) \text{ to the holomorphic space}$ 
6: end if
7:  $\forall i \in 1..|\mathcal{A}|, \Delta\theta_i := \text{Arg}(\Gamma(p_{i+1})/\Gamma(p_i))$ 
8:  $Z_0 := \{z_0 \in \Omega \mid \Gamma(z_0) = 0\}$ 
9: for each  $z_0$  in  $Z_0$  do
10:   $i := \text{index of the closest point in } \mathcal{A} \text{ to } z_0$ 
11:   $\Delta\theta_i := \Delta\theta_i - 2\pi$ 
12: end for
13:  $\theta_1 = \text{Arg}(\Gamma(p_1))$ 
14:  $\forall i \in 2..|\mathcal{A}|, \theta_i := \theta_{i-1} + \Delta\theta_i$ 

```

6 and 7 show some results.

7 Results

We implemented the three projection algorithms that are presented in this paper using Matlab. We used the CVX 2.1 software for Disciplined Convex Programming (DCP) [Grant et al. 2008] to model these optimization problems. CVX automatically converts a convex problem that adheres to the DCP rules to a form that can then be solved by a wide range of commercial solvers. We used Mosek [ApS 2015] as the underlying solver. The running time of each projection algorithm is dominated by that of the particular convex optimization (13),(22),(28),(29). On a problem with moderate size with 50 Cauchy basis functions and 300 boundary samples, the solve took approximately 1 second, and we observed that the running time roughly scales linearly with the problem size. Solving (22) is somewhat slower than (13),(28),(29) as CVX 2.1 does not

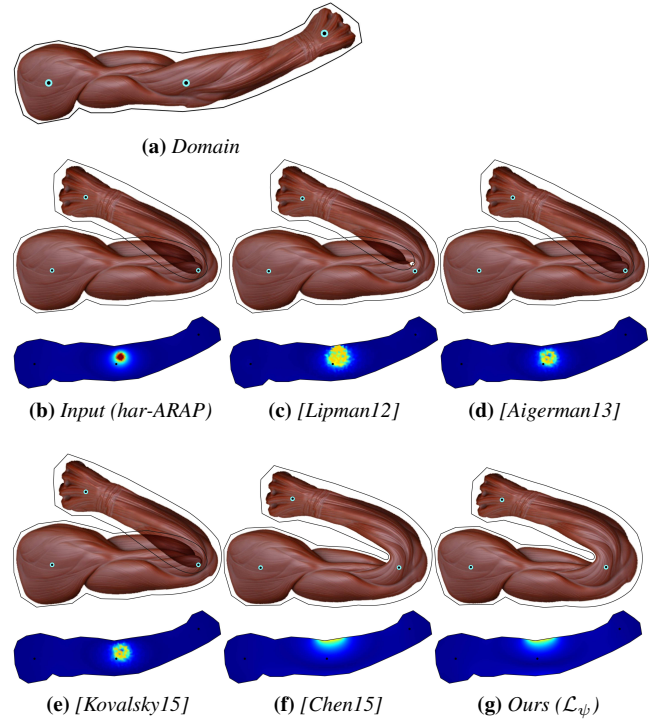


Figure 10: Arm with soft positional constraints. All the methods satisfy the positional constraints. Our result resembles the one of Chen’s and is smoother compared to Lipman’s. The color plots visualize the isometric distortion τ_f .

have built-in support for exponential cones. Nonetheless, they all have comparable running times. On our largest model with 150 basis functions and $|\mathcal{A}| = 1000$ samples, the running time was 3 seconds.

We used the same default parameters for all the figures as explained next, unless stated otherwise. The lower bound σ was set to 0.2, and the upper bound Σ was set to 5. The upper bound on the conformal distortion k was set to 0.7. We used the following definition for the isometric distortion $\tau_f(z) = \max(\Sigma_f(z), 1/\sigma_f(z))$ and visualized $\tau_f(z)$ (or $k_f(z)$ if explicitly noted) on the source domain Ω using Matlab’s jet color scheme (red for high values and blue for low). Wherever we used positional constraints (for the \mathcal{L}_ψ projection), we marked the target position of the point with a cyan disk, while the actual deformed position is marked with a smaller black disk to visually illustrate satisfaction of the constraints.

The following unconstrained methods were used to generate the input mappings for our experiments: As-Rigid-As-Possible (ARAP) shape deformation [Sorkine and Alexa 2007]; unconstrained harmonic ARAP (har-ARAP) which is a simplified version of [Chen and Weber 2015] that minimizes the same ARAP energy within the same harmonic subspace, but without the bounded distortion constraints; Cauchy coordinates [Weber et al. 2009] with or without positional constraints; and harmonic coordinates (HC) [Joshi et al. 2007]. Unconstrained methods tend to concentrate distortion at strategic locations. For example, [Weber et al. 2009] generate results which are holomorphic but typically not conformal, having vanishing derivative at several points inside the domain. This is illustrated in Figures 5,12. The zeros of the derivatives are easily spotted by looking at the color visualization for τ_f since at these zeros, σ vanishes and τ_f becomes infinite (indicated by red regions). We compared our results with the projection methods of [Aigerman and Lipman 2013; Kovalsky et al. 2015] as well as with the

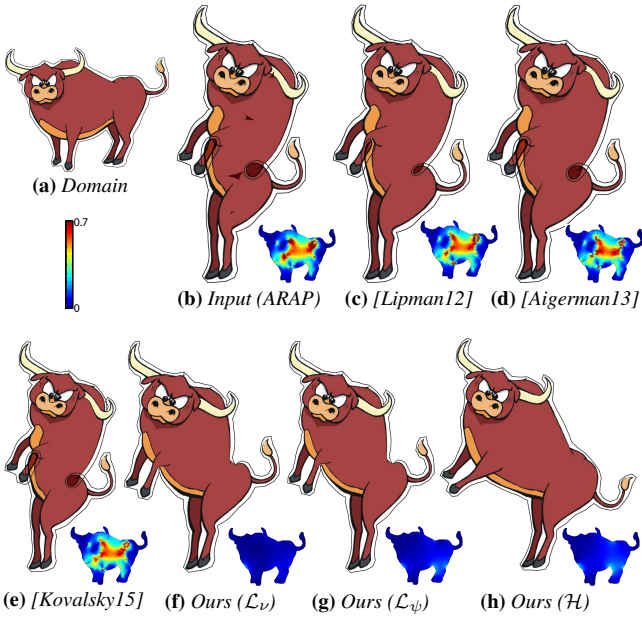


Figure 11: Bull. Projecting a reference mapping that was produced with the ARAP method. Comparison of the results of our three spaces with alternative methods. The color visualization shows the conformal distortion k_f .

constrained optimization methods of [Lipman 2012] and its smooth harmonic counterpart [Chen and Weber 2015]. This is done by minimizing the following quadratic energy subject to the convexified conformal and isometric constraints

$$\min_f |f_z - g_z|^2 + |f_{\bar{z}} - g_{\bar{z}}|^2 + \lambda_{p2p} |f(c_i) - f_i|^2 \quad (32)$$

[Aigerman and Lipman 2013] typically required 4-8 iterations to converge. The running time of each iteration is dominated by the solution of the underlying quadratic program (QP) and is similar to ours on small to moderate size meshes. [Kovalsky et al. 2015] typically required 30-70 iterations, where each iteration requires only a linear solve (with prefactored left hand side matrix) and generally was faster. Nevertheless, the acceleration procedure proposed by Kovalsky et al. is only applicable to bound conformal distortion. For appropriate comparison, we incorporated isometric constraints similarly to Aigerman’s (QP). This resulted in slower computation times. Nonetheless, the purpose of doing so is to assess the qualitative behavior of the algorithm only. We allowed the methods to perform up to 300 iterations before declaring failure of convergence. For [Lipman 2012; Chen and Weber 2015], 25 iterations were performed, which were enough for convergence. The running time of each iteration is dominated by solving an SOCP and is similar to ours, though our method performs a single iteration. Moreover, in the mesh-based methods, the number of variables and inequality constraints is linear in the number of mesh vertices, whereas in our method, these depend solely on the complexity of the boundary. Hence, if one wants to use the mesh-based methods to produce results which are much smoother, it would take significantly longer to compute. For all these experiments, we used meshes with roughly 10,000 triangles.

Both [Lipman 2012] and [Chen and Weber 2015] require proper initialization to operate. These are encoded with local frames that determine the convexified subspace. For the method to *guarantee* feasibility and to produce locally injective mappings, the frames must be extracted from a bounded distortion mapping. Hence, in all the results presented, we used the identity mapping for initial-

ization (which is the only available choice). In many cases where g is significantly different from the identity mapping, the projected mapping was quite far from the expected result even after many iterations (the nonconvex minimization is essentially trapped in a local minimum). This is evident in Figures 1,14,15. We also tried initialization with frames extracted from the input g . However, these two methods have the property that they preserve the zeros set of g_z . Hence, in cases where the reference mapping contains singularities, the result will *not* be locally injective (since $g_z \neq 0$ is a necessary condition for local injectivity). We refer to [Aigerman et al. 2014, Section 3.3] for further discussion. The results obtained with this initialization visually resembled those of [Aigerman and Lipman 2013] and [Kovalsky et al. 2015], and were omitted. Hence, all the included results for [Lipman 2012] and [Chen and Weber 2015] are locally injective (away from the boundary) and obey the bounds.

Let us consider the example of the pants in Figure 12. Figure 13 shows the steps that [Lipman 2012] takes and the evolution of the rotation which is unintuitive and result in an unnecessary small loop on the boundary. [Chen and Weber 2015] took similar steps. The difference between Chen’s result and Lipman’s result is that Chen’s favored smoothness over fidelity to the input due to the restriction to the harmonic space.

The \mathcal{L}_ν and \mathcal{L}_ψ spaces address the rotations problem specifically and without iterations. In Figure 1, where the input mapping is a spiral, the wrong rotations that were obtained by [Lipman 2012] and [Chen and Weber 2015] are clearly illustrated. The same behavior is illustrated on a large collection of deformations with significant bending of a bar model in Figure 7. The superiority of our method with respect to the correct handling of rotations is also illustrated in other examples; see Figures 2, 6, 8, and 14.

Figures 9 and 10 illustrate the effect of smoothness. Lipman’s result is valid and better resembles the reference mapping compared to our result. However, ours (as well as Chen’s which is remarkably similar) is much smoother and have lower average distortion.

Figure 11 compares our three spaces with each other and with the alternative methods. As evident, \mathcal{L}_ψ and \mathcal{L}_ν produce quite similar results. Such a behavior was generally observed when no positional constraints are specified. The result of the \mathcal{H} space is visually more distant from the input (especially near the tail which was rotated by approximately 45 degrees). Figures 8 and 9 use the same reference mapping as input and allow for comparison of the result of \mathcal{L}_ν with that of \mathcal{L}_ψ when positional constraints are used. A similar comparison between \mathcal{L}_ν and \mathcal{L}_ψ is available in Figures 14 and 15.

In all our experiments, [Aigerman and Lipman 2013] and [Kovalsky et al. 2015] performed quite similar to each other and tend to produce mappings which are not locally injective in the presence of vanishing points of g_z . It seems that since there is nothing specific in these methods that attempts to target local injectivity at mesh *vertices*, they tend to favor (energetically) singularities of the derivative if these are present in g . This is evident in all the results except in Figure 4 in which local injectivity was obtained. The lack of local injectivity can sometimes be spotted visually in the results by looking at the black curve that surrounds the images. While the curve is allowed to self-intersect, a curve with “loops” whose turning number is different than 1 [Weber and Zorin 2014] implies the lack of local injectivity.

Finally, we provide additional results in a separate supplementary material document.




Figure 12: *Pants.* (g) Source domain. (h) visualization of the isometric distortion. (c,d) Not locally injective. (b,e) Rotated to the wrong direction.




Figure 13: *The first 5 iterations of [Lipman12] corresponding to Figure 12(b). Local frames are initialized with identity.*

8 Summary and Discussion

We presented a novel framework for the efficient computation of operators that take an arbitrary planar mapping as input and are *guaranteed* to produce in return a similar high quality mapping that adheres to a set of very strict demands. They should be C^∞ , locally injective, and their geometric distortion must be bounded according to some user specified bounds.

We achieve this by exploiting unique mathematical properties of complex holomorphic functions and by a sophisticated change of variables, which allows us to characterize the nonconvex space of bounded distortion harmonic mappings using three spaces. The transformation to and from these spaces can be easily computed using simple operators and the underlying optimization problems can be solved efficiently in an optimal manner using off-the-shelf solvers.

The first convex space \mathcal{H} is shown to be in full correspondence with \mathcal{BD} (the space of bounded distortion harmonic mappings). This is somewhat surprising since \mathcal{BD} is highly nonconvex. Theoretically, this is a fascinating outcome, however, the transition to the convex space has its price. Typical objective functions, such as Equation (32) and positional constraints, become nonconvex in the new variables. In order to enjoy the benefits of convex programming, we designed an objective (10) that is convex but does not explicitly

capture the rotational part of the Jacobian. As such, the projected mapping does not always provide good enough visual fidelity to the input. We believe that the full strength of the \mathcal{H} space will be utilized in other applications such as parameterization, where position and orientation control are less vital.

The second space, \mathcal{L}_ν , is based on a logarithmic characterization of the conformal part of harmonic mappings. In contrast to \mathcal{H} , it controls orientation explicitly. This space is not convex, yet the detailed analysis in Section 6.1 shows that it is nearly so. To benefit from convex programming, we use a convexification which does not depend on user defined parameters and is proven to be feasible.

The last space, \mathcal{L}_ψ , is specifically designed to address the inability of the \mathcal{H} and \mathcal{L}_ν spaces to incorporate positional constraints. Similar to \mathcal{L}_ν it is based on a logarithmic characterization and is not convex. However, in Section 6.2 we show that the space of all conformal mappings with bounded distortion is a (nonempty) convex subspace of \mathcal{L}_ψ . Our strategy is then to first solve a convex program for the “closest” conformal mapping (without positional constraints). We then show that by fixing the conformal part of the obtained mapping (and setting the anti-holomorphic part $\bar{\Psi}$ free) the problem of finding the closest mapping with soft positional constraints is again a convex program. Such a two-step solution is not optimal but it performs well in practice. If positional constraints are not required, \mathcal{L}_ν would be a wiser choice to use.

Limitations. Our method has two main limitations. First, similar to [Aigerman and Lipman 2013; Kovalsky et al. 2015], our energy targets similarity to a given reference mapping rather than an independent deformation functional such as low *average* distortion. Thus, in its current form, it cannot be used as a standalone deformation tool but rather as a complement to (any) existing methods.

The second limitation is the inability to *guarantee* satisfaction of positional constraints. While none of the currently available methods has such ability, it is still highly desirable. One of the problems that arises is that with such constraints even the full space of bounded distortion mappings might be empty.

Future work. The immediate avenue for extension we would like to pursue is the application of our spaces to curved surfaces, which is very useful for parameterization and quadrangulation applications. While this seems to be possible, it would require a different discretization and it is not fully clear yet whether the smooth properties of holomorphic functions would carry over to the discrete setting. Finally, the existence of a convex space such as \mathcal{H} that completely characterizes the space of bounded distortion harmonic mappings raises a question about the existence of other convex spaces with complete correspondence that have alternative properties and can be employed in different scenarios, allowing the addition of other types of constraints and design of other objective functions in a convex framework.

Acknowledgements

This research was partially funded by the Israel Science Foundation (grants No. 1869/15 and 2102/15). We thank Edward Chien for insightful discussions and proofreading the paper, and the anonymous reviewers for their valuable comments and suggestions.

References

- AHLFORS, L. 1979. *Complex analysis*, vol. 7. McGraw-Hill Education.
- AIGERMAN, N., AND LIPMAN, Y. 2013. Injective and bounded distortion mappings in 3D. *TOG* 32, 4, 106.




Figure 14: Orange bar with soft positional constraints. Due to the symmetric structure of the positional constraints (b,e) were unable to deviate much from the identity mapping that was used to initialize them and failed to satisfy the positional constraints.




Figure 15: Orange bar without positional constraints. Our method as well as [Chen and Weber 2015], which uses the same harmonic subspace, produced smooth results with overall low average isometric distortion while the other methods tend to concentrate the distortion near the zeros of g_z (red areas in (a)).

- AIGERMAN, N., PORANNE, R., AND LIPMAN, Y. 2014. Lifted bijections for low distortion surface mappings. *TOG* 33, 4, 69.
- ALEXA, M. 2002. Recent advances in mesh morphing. In *Computer Graphics Forum*, vol. 21, 173–198.
- APS, M. 2015. *The MOSEK optimization toolbox for MATLAB manual. Version 7.1 (Revision 28)*.
- BELL, S. R. 1992. *The Cauchy transform, potential theory and conformal mapping*, vol. 7. CRC press.
- BOMMES, D., LÉVY, B., PIETRONI, N., PUPPO, E., SILVA, C., TARINI, M., AND ZORIN, D. 2013. Quad-mesh generation and processing: A survey. *Computer Graphics Forum* 32, 6, 51–76.
- BOTSCH, M., AND SORKINE, O. 2008. On linear variational surface deformation methods. *TVCG* 14, 1, 213–230.
- BOYD, S., AND VANDENBERGHE, L. 2004. *Convex optimization*. Cambridge University Press.
- CHEN, R., AND WEBER, O. 2015. Bounded distortion harmonic mappings in the plane. *ACM TOG* 34, 4, 73.
- CHIEN, E., CHEN, R., AND WEBER, O. 2016. Bounded distortion harmonic shape interpolation. *ACM TOG* 35, 4.
- DUREN, P. 2004. *Harmonic mappings in the plane*. Cambridge University Press.
- FLOATER, M., AND HORMANN, K. 2005. Surface Parameterization: a Tutorial and Survey. *Adv. Multi. Geom. Modelling*.
- FLOATER, M. 1997. Parametrization and smooth approximation of surface triangulations. *CAGD* 14, 3, 231–250.
- GRANT, M., BOYD, S., AND YE, Y., 2008. Cvx: Matlab software for disciplined convex programming.
- HORMANN, K., AND GREINER, G. 2000. Mips: An efficient global parametrization method. *Curve and Surface Design: Saint-Malo 99*, 153–162.
- JOSHI, P., MEYER, M., DEROSE, T., AND GREEN, B. 2007. Harmonic coordinates for character articulation. *ACM TOG* 26, 3, 71.
- KHAREVYCH, L., SPRINGBORN, B., AND SCHRÖDER, P. 2006. Discrete conformal mappings via circle patterns. *ACM Trans. Graph.* 25, 2, 412–438.
- KOVALSKY, S. Z., AIGERMAN, N., BASRI, R., AND LIPMAN, Y. 2014. Controlling singular values with semidefinite programming. *ACM TOG* 33, 4.
- KOVALSKY, S. Z., AIGERMAN, N., BASRI, R., AND LIPMAN, Y. 2015. Large-scale bounded distortion mappings. *ACM TOG* 34, 6, 191.
- LEVI, Z., AND ZORIN, D. 2014. Strict minimizers for geometric optimization. *ACM TOG* 33, 6, 185.
- LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. Least squares conformal maps for automatic texture atlas generation. *ACM TOG* 21, 3, 362–371.
- LIPMAN, Y. 2012. Bounded distortion mapping spaces for triangular meshes. *ACM TOG* 31, 4, 108.
- LIU, L., ZHANG, L., XU, Y., GOTSMAN, C., AND GORTLER, S. 2008. A local/global approach to mesh parameterization. *Computer Graphics Forum* 27, 5, 1495–1504.
- PORANNE, R., AND LIPMAN, Y. 2014. Provably good planar mappings. *ACM TOG* 33, 4, 76.
- SCHÜLLER, C., KAVAN, L., PANOZZO, D., AND SORKINE-HORNUNG, O. 2013. Locally injective mappings. In *Computer Graphics Forum*, vol. 32, 125–135.
- SHEFFER, A., PRAUN, E., AND ROSE, K. 2006. Mesh Parameterization Methods and Their Applications. *Found. Trends. Comput. Graph. Vis.* 2, 2, 105–171.
- SORKINE, O., AND ALEXA, M. 2007. As-rigid-as-possible surface modeling. In *SGP*, 109–116.
- SORKINE, O. 2006. Differential representations for mesh processing. *Computer Graphics Forum* 25, 4, 789–807.
- SPRINGBORN, B., SCHRÖDER, P., AND PINKALL, U. 2008. Conformal equivalence of triangle meshes. *ACM TOG* 27, 3, 77.
- TUTTE, W. 1963. How to draw a graph. *Proc. London Math. Soc* 13, 3, 743–768.
- WEBER, O., AND GOTSMAN, C. 2010. Controllable conformal maps for shape deformation and interpolation. *ACM TOG* 29, 4.
- WEBER, O., AND ZORIN, D. 2014. Locally injective parametrization with arbitrary fixed boundaries. *TOG* 33, 4, 75.

WEBER, O., BEN-CHEN, M., AND GOTSMAN, C. 2009. Complex barycentric coordinates with applications to planar shape deformation. *Computer Graphics Forum* 28, 2, 587–597.

WEBER, O., MYLES, A., AND ZORIN, D. 2012. Computing extremal quasiconformal maps. *Computer Graphics Forum* 31, 5, 1679–1689.

WEBER, O. 2010. *Hybrid Methods for Interactive Shape Manipulation*. PhD thesis, Technion - Israel Institute of Technology.

WOLBERG, G. 1998. Image morphing: a survey. *The visual computer* 14, 8, 360–372.

A One-to-one correspondence - Proof of Theorem 7

Recalling Theorem 7:

$\mathcal{F} : \mathcal{BD} \rightarrow \mathcal{H}$ is a one-to-one correspondence (bijection).

The proof of the theorem is based on four lemmas.

Lemma 11. *Let Φ_1, Φ_2 be conformal mappings on simply-connected domain $\Omega \subset \mathbb{C}$. If $\forall w \in \partial\Omega$, $|\Phi_1'(w)| = |\Phi_2'(w)|$, then $\forall z \in \Omega$, $\Phi_1'(z) = e^{i\alpha}\Phi_2'(z)$ for some constant $\alpha \in \mathbb{R}$.*

Proof. Since Φ_1, Φ_2 are conformal, their derivatives Φ_1', Φ_2' are nonvanishing holomorphic functions. Hence, there exists a holomorphic branch for $\log(\Phi_1'(z))$ and $\log(\Phi_2'(z))$. Furthermore, $h_1 = \text{Re}(\log(\Phi_1'(z))) = \ln(|\Phi_1'(z)|)$ and $h_2 = \text{Re}(\log(\Phi_2'(z))) = \ln(|\Phi_2'(z)|)$ are harmonic real-valued functions, and $\text{Im}(\log(\Phi_1'(z)))$, and $\text{Im}(\log(\Phi_2'(z)))$ are their harmonic conjugates respectively. Since we know that $|\Phi_1'(w)| = |\Phi_2'(w)|$, we have that h_1, h_2 agree on the boundary and since the solution to the Dirichlet problem is unique we must have

$$\ln(|\Phi_1'(z)|) = \ln(|\Phi_2'(z)|), \quad (33)$$

everywhere. Since $\text{Im}(\log(\Phi_1'(z)))$ and $\text{Im}(\log(\Phi_2'(z)))$ are harmonic conjugates of the same function, they must be identical up to some real constant α

$$\text{Im}(\log(\Phi_1'(z))) = \alpha + \text{Im}(\log(\Phi_2'(z))). \quad (34)$$

Multiplying (34) by i and summing with (33) gives

$$\log(\Phi_1'(z)) = \log(\Phi_2'(z)) + i\alpha. \quad (35)$$

Finally, exponentiating both sides concludes the proof. \square

Lemma 12. *If f is a locally injective planar harmonic mapping, then Φ in the decomposition $f = \Phi + \bar{\Psi}$ (Equation (8)) is a conformal mapping.*

Proof. Since Φ is holomorphic we only need to show that Φ' does not vanish. f is locally injective and $f_z = \Phi'$ hence $|f_z| = |\Phi'| > |f_{\bar{z}}| \geq 0$ (Equation (4)). \square

Lemma 13. $\mathcal{F} : \mathcal{BD} \rightarrow \mathcal{H}$ is injective.

Proof. Given $f_1, f_2 \in \mathcal{BD}$ such that $\mathcal{F}(f_1) = \mathcal{F}(f_2)$ we need to show that $f_1 = f_2$. Using Equation (8) we can write

$$f_1(z) = \Phi_1(z) + \bar{\Psi}_1(z), \quad \Psi_1(z_0) = 0, \quad (36a)$$

$$f_2(z) = \Phi_2(z) + \bar{\Psi}_2(z), \quad \Psi_2(z_0) = 0, \quad (36b)$$

Since $\mathcal{F}(f_1) = \mathcal{F}(f_2)$ and based on Definition 5 we have

$$\Psi_1(z) = \Psi_2(z), \quad \forall z \in \Omega, \quad (37a)$$

$$|\Phi_1'(w)| = |\Phi_2'(w)|, \quad \forall w \in \partial\Omega, \quad (37b)$$

$$\text{Arg}(\Phi_1'(z_0)) = \text{Arg}(\Phi_2'(z_0)), \quad (37c)$$

$$f_1(z_0) = f_2(z_0). \quad (37d)$$

Since f_1, f_2 are locally injective, Lemma 12 asserts that Φ_1, Φ_2 are conformal. Together with Equation (37b), the conditions for Lemma 11 hold and we can write $\Phi_1'(z) = e^{i\alpha}\Phi_2'(z)$. This means that $\Phi_1'(z), \Phi_2'(z)$ differ by a global rotation of angle α . From (37c) we must have that $\alpha = 2\pi n$, for some integer n , hence $e^{i\alpha} = 1$ and we get that $\Phi_1'(z) = \Phi_2'(z)$. Furthermore, the antiderivatives $\Phi_1(z), \Phi_2(z)$ must be identical up to a complex constant. Finally, by evaluating Equation (36) at z_0 and using Equation (37d), we obtain $\Phi_1(z) = \Phi_2(z)$, which together with Equation (37a) lead to $f_1(z) = f_2(z)$. \square

Lemma 14. $\mathcal{F}^{-1} : \mathcal{H} \rightarrow \mathcal{BD}$ is a right inverse of \mathcal{F} .

Proof. We need to prove that for any $h = \{\Psi(z), r(w), d, \theta\} \in \mathcal{H}$, $\mathcal{F}(\mathcal{F}^{-1}(h)) = h$.

Let us first show that $\mathcal{F}^{-1}(h)$ is well-defined for any $h \in \mathcal{H}$, and that it is unique and produces a mapping f which is in \mathcal{BD} . From Condition (7c) we have $|\Psi'(w)| \leq r(w) - \sigma$. Since σ is strictly positive, $r(w)$ is also strictly positive, and $\ln(r(w))$ is a well defined real (single-valued) continuous function on $\partial\Omega$. This means that we can uniquely solve the Dirichlet problem, obtaining the harmonic function $\xi(z)$. Moreover, since Ω is simply-connected, the Hilbert transform of $\xi(z)$ exists and it is unique up to a constant, hence the function $l(z) = \xi(z) + i\zeta(z)$ is holomorphic and is unique for the choice $\zeta(z_0) = \theta$. The function $e^{l(z)}$ is also holomorphic and so is its antiderivative $\Phi = \int e^l$. For uniqueness, we set $\Phi(z_0) = d$. $f = \Phi + \bar{\Psi}$ is a sum of holomorphic and anti-holomorphic functions, hence it is clearly harmonic and it is left to show that $f \in \mathcal{BD}$.

Differentiating f with respect to z gives $f_z(z) = \Phi'(z) = e^{\xi(z)}e^{i\zeta(z)}$, therefore on the boundary we have $|f_z(w)| = |\Phi'(w)| = |e^{\xi(w)}e^{i\zeta(w)}| = e^{\xi(w)} = r(w)$. Similarly, by differentiating with respect to \bar{z} we have $|f_{\bar{z}}(w)| = |\Psi'(w)|$, and with that we can use (7) to determine that the conditions in (6) hold on the boundary of the domain. In order to show that $f \in \mathcal{BD}$, we need to show that the conditions in (6) hold everywhere. To this end we apply Theorem 4 of [Chen and Weber 2015] which asserts that it is sufficient to bound the distortion of a harmonic mapping on its boundary to imply global bounds as long as $f_z(z)$ does not vanish in the interior. Luckily, the latter condition holds since $f_z(z)$ is the exponent of $l(z)$. We have shown that $\mathcal{F}^{-1}(h) = f \in \mathcal{BD}$.

It remains to show that $\mathcal{F}(f) = h$. The harmonic decomposition is unique for the choice $\Psi(z_0) = 0$. Moreover, for this choice we have $f(z_0) = \Phi(z_0) = d$. Lastly, $\text{Arg}(\Phi'(z_0)) = \zeta(z_0) = \theta$. \square

Finally, a corollary of Lemma 14 is that \mathcal{F} is surjective, and since it is also injective (Lemma 13) it is a bijection which concludes the proof of Theorem 7. Moreover, since \mathcal{F} is bijective, \mathcal{F}^{-1} is also a left inverse of \mathcal{F} .

B Feasibility of the convex \mathcal{L}_ν space

Proof. First, let ν be zero (this implies that the mapping is conformal). Condition (20a) is immediately satisfied. Substituting $\nu = 0$ in (20b) and (20c), taking \ln of (20c) and further rearranging gives

$$\ln \sigma \leq \text{Re}(l(w)) \leq \ln \Sigma \quad \forall w \in \partial\Omega. \quad (38a)$$

Since $\ln \sigma \leq \ln \Sigma$, we can (for example) choose $l(w)$ to be constant (holomorphic) function $l(w) = \frac{1}{2}(\ln \sigma + \ln \Sigma) + iC$ for any real constant C , which clearly satisfy (38a). \square

On the Convexity and Feasibility of the Bounded Distortion Harmonic Mapping Problem - Supplementary material

Zohar Levi
Victoria University of Wellington, New Zealand

Ofir Weber
Bar Ilan University, Israel

Supplementary Material

Below we provide some additional results of our method as well as comparisons with other methods.




Figure 1: Projection of a simple holomorphic mapping. (a) The image of the input mapping, $f(z) = (z - 0.05)^2$ applied to the unit disk domain (g). The mapping is not locally injective and has a singular point at $z = 0.05$ where $f_z = 2(z - 0.05)$ vanishes. The isometric distortion at the singularity is infinite since $\sigma_f(0.05) = 0$. This is evident by the color visualization of the isometric distortion (h). Our projection operator removes the singularity while [Aigerman13] and [Kovalsky15] maintains it. [Lipman12] pushes the singularity to the boundary and similarly [Chen15] whose result is smoother.

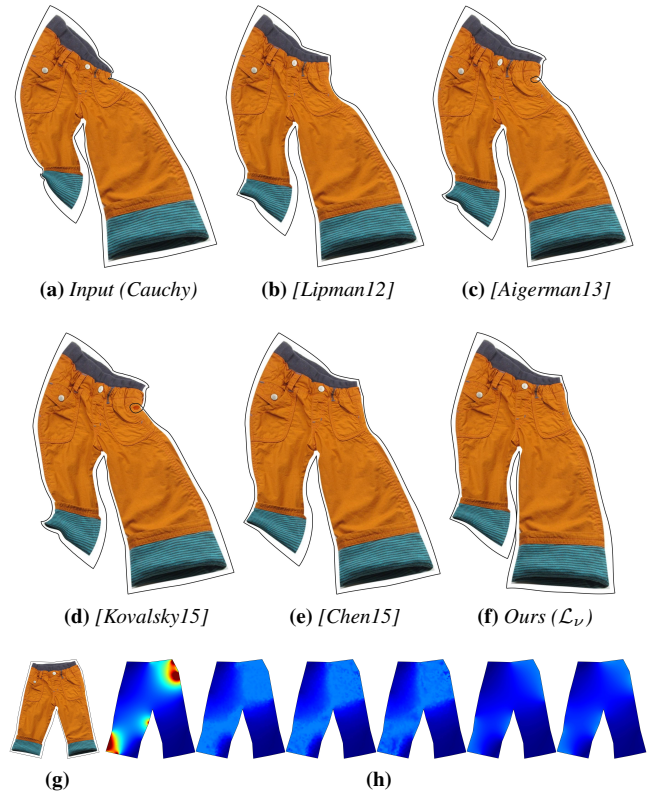


Figure 2: Pants. Constraints: $\sigma = 0.4$, $\Sigma = 2.5$. (g) Domain. (h) Isometric distortion τ_f visualization of (a-f).

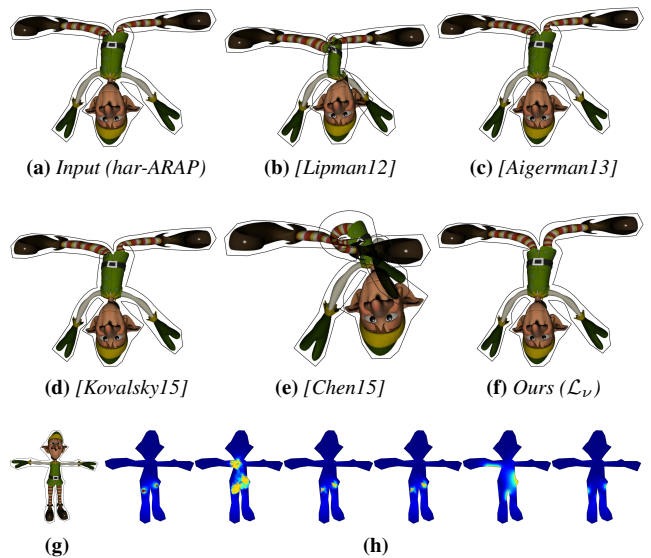


Figure 3: Elf. (g) Domain. (h) Isometric distortion τ_f visualization of (a-f).




Figure 4: Pants. (g) Domain. (h) Isometric distortion τ_f visualization of (a-f).




Figure 5: Square. (g) Domain. (h) Isometric distortion τ_f visualization of (a-f).




Figure 6: Troll. (g) Domain. (h) Visualization of the conformal distortion k_f .




Figure 7: Taz. (g) Domain. (h) Isometric distortion τ_f visualization of (a-f).