

GPU-Accelerated Locally Injective Shape Deformation

RENJIE CHEN, Max Planck Institute for Informatics, Germany
OFIR WEBER, Bar-Ilan University, Israel



Fig. 1. Locally injective isometric deformation of Rex, a doubly-connected planar domain (left). Note the hole between the legs. Our harmonic subspace has 392 DOF. Our solver converges after 6 iterations to the result depicted, which is computed in 0.072s when implemented on the GPU, and 0.375s on the CPU. We compare against an unmodified Newton solver which halts prematurely after 9 iterations at a high energy configuration due to a non-positive definite Hessian. The Newton-Eigen solver performs a full eigen-factorization Hessian modification at each iteration. It converges after 32 iterations and takes 2.467s which is $\times 34$ times slower than our GPU solver. The L-BFGS solver performs 1,818 iterations and is not competitive in terms of runtime.

We present a highly efficient planar meshless shape deformation algorithm. Our method is based on an unconstrained minimization of isometric energies, and is guaranteed to produce C^∞ locally injective maps by operating within a reduced dimensional subspace of harmonic maps. We extend the harmonic subspace of [Chen and Weber 2015] to support multiply-connected domains, and further provide a generalization of the bounded distortion theorem that appeared in that paper. Our harmonic map, as well as the gradient and the Hessian of our isometric energies possess closed-form expressions. A key result is a simple-and-fast *analytic* modification of the Hessian of the energy such that it is positive definite, which is crucial for the successful operation of a Newton solver. The method is straightforward to implement and is specifically designed to harness the processing power of modern graphics hardware. Our modified Newton iterations are shown to be extremely effective, leading to fast convergence after a handful of iterations, while each iteration is fast due to a combination of a number of factors, such as the smoothness and the low dimensionality of the subspace, the closed-form expressions for the differentials, and the avoidance of expensive strategies to ensure positive definiteness. The entire pipeline is carried out on the GPU, leading to deformations that are significantly faster to compute than the state-of-the-art.

CCS Concepts: • **Computing methodologies** \rightarrow *Computer graphics; Animation; Image manipulation; Shape analysis*;

This research was partially funded by the Israel Science Foundation (grants No. 1869/15 and 2102/15) and by the Max Planck Center for Visual Computing and Communication.

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 republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 Association for Computing Machinery.
0730-0301/2017/11-ART214 \$15.00

<https://doi.org/10.1145/3130800.3130843>

Additional Key Words and Phrases: injective maps, harmonic maps, shape deformation, Newton method, GPU

ACM Reference Format:

Renjie Chen and Ofir Weber. 2017. GPU-Accelerated Locally Injective Shape Deformation. *ACM Trans. Graph.* 36, 6, Article 214 (November 2017), 13 pages. <https://doi.org/10.1145/3130800.3130843>

1 INTRODUCTION

Shape deformation is a classical problem in computer graphics and animation that comes in a variety of forms and settings. Essential to all variants is the requirement to compute a map between different spaces with some desirable properties. Common choices include maps between curved manifolds embedded in \mathbb{R}^3 , volumetric maps between solid shapes in \mathbb{R}^3 , and mixed dimensional maps where, for example, a curved 2-manifold embedded in \mathbb{R}^3 is mapped to \mathbb{R}^2 (a parameterization). In this work we deal with the problem of *planar* maps where both the source and target domains are subsets of \mathbb{R}^2 . Due to its relative simplicity, this setting is the most fundamental one that serves as a bridge for all other developments and generalizations, and as such it attracts a large amount of research efforts since the early days of graphics. Depending on the precise user requirements, even this simplified setting is challenging to deal with.

Some requirements such as smoothness are easily achievable while other, such as local injectivity or strict bound on the induced metric distortion are hard (sometimes impossible) to guarantee. In this work we consider an interactive deformation system that is driven by user specified positional constraints. That is, the user selects a small set of points in the source domain which is visualized as a textured image, and interactively drags them around to new

positions. The goal is to be able to compute a map that adheres to these positional constraints yet at the same time is visually plausible. Plausibility is usually quantified by smoothness and the amount of metric distortion, which can be either minimized on the average, or bounded above by some user-defined amount. Moreover, the map is expected to be locally injective, preventing degeneracies of the image and local overlaps while allowing global ones. These requirements pose a great computational challenge since they usually lead to a nonconvex optimization problem with many degrees of freedom for which finding an optimal solution (or even just a feasible one) is often intractable. Recent advances in the computation of locally injective and/or bounded distortion maps allow artists to produce high quality results with greater than ever success rates but the performance of current methods is still significantly worse compared to simple linear methods [Weber et al. 2009].

We cast the deformation problem as a nonlinear nonconvex unconstrained minimization problem that is performed in a reduced deformable subspace of harmonic maps, and allow the user to pick among a variety of smooth energies that measure isometric distortion. Harmonic functions have many appealing properties and are widely used in computer graphics applications. In particular, holomorphic and harmonic maps are utilized extensively in the context of planar shape deformation and interpolation [Chen and Weber 2015; Chien et al. 2016a; Hefetz et al. 2017; Levi and Weber 2016; Weber 2010, 2017; Weber et al. 2009, 2011; Weber and Gotsman 2010]. However, this entire line of work was restricted to simply-connected domains (topological disks). Our first contribution in this work is broadening the scope of previous research on harmonic maps from simply-connected to multiply-connected domains (disks with finite number of holes) such as those appearing in Figures 1, 3, 5. In particular, in Section 4 we extend the harmonic subspace of the Bounded Distortion Harmonic Mappings (BDHM) method [Chen and Weber 2015] to multiply-connected domains, as well as generalize the bounded distortion theorem of BDHM (Section 4.1). Furthermore, in Section 8 we provide a much tighter analysis of the Lipschitz constants provided in BDHM, leading to more accurate validation of the map and faster convergence.

We construct a custom-made Newton solver with unique features to efficiently address the deformation problem. The quality of the produced maps is superior and comparable to the other methods that operate within the harmonic subspace. However, our method is orders of magnitude faster than existing techniques. Our optimization is based on a second-order Newton approach which ensures that the convergence rate is high. In general, a single Newton iteration can be slow due to various reasons such as the need to approximate the Hessian numerically using finite differences (when analytical expressions are unavailable), as well as the need to modify the Hessian such that it becomes positive definite (which is necessary to ensure convergence to a local minimum).

Our novel solver overcomes these issues due to several unique properties, such as having closed-form expressions for the map, and its differentials, and an analytic modification scheme for the Hessian to ensure its positive definiteness. Combining these benefits leads to a Newton iteration that is considerably faster than a standard Newton iteration. Together with the small number of iterations that are typically required for convergence, the total running time

of our solver is significantly shorter compared to state-of-the-art solvers. Moreover, every iteration of our solver is embarrassingly parallelized and perfectly suits implementation on massively parallel graphics processing units due to the structure and regularity of the problem. This allows for unprecedented performance.

2 PREVIOUS WORK

Linear blend skinning [Kavan et al. 2007] and methods that are based on barycentric coordinates [Hormann and Floater 2006; Lipman et al. 2007; Weber et al. 2009] use precomputed basis functions [Jacobson et al. 2011] and formulate the deformation as a linear combination of these basis functions with some coefficients that are automatically deduced from the user input. For instance a control polygon/mesh or rotations that are extracted from a skeletal structure. Despite their inability to produce maps with good distortion control, these methods are still highly popular due to their speed and in particular their suitability for GPU implementation.

More geometric motivated linear approaches merely require solving a single linear system with variable right hand sides where the left hand side matrix is fixed and can be factored in preprocessing [Igarashi et al. 2005; Lipman et al. 2005; Weber et al. 2009, 2007; Zayer et al. 2005]. Under extreme deformations, artifacts arise which lead to the introduction of more sophisticated nonlinear approaches.




Fig. 2. Comparison with BDHM using the same harmonic subspace. We set user distortion bounds for BDHM: $(k, \sigma_1, \sigma_2) = (0.8, 5, 0.3)$. The resulted map is bounded distortion with bounds $(0.50, 4.3, 0.31)$. Without explicit distortion bounds, our method minimizes the symmetric Dirichlet energy producing a map with bounds $(0.32, 2.25, 0.43)$. Our GPU solver converges after 8 iterations with total runtime of 0.0297s. In comparison, BDHM requires 28 iterations of second order cone programming [ApS 2017], with total runtime 3.71s which is $\times 125$ slower.

Recently, a successful line of approaches have been designed to specifically address quality concerns with a priority on producing injective maps [Bright et al. 2017; Fu et al. 2015; Schüller et al. 2013; Smith and Schaefer 2015] with bounded amount of distortion [Aigerman and Lipman 2013; Aigerman et al. 2014; Chen and Weber 2015; Chien et al. 2016b; Kovalsky et al. 2015; Lipman 2012; Poranne and Lipman 2014]. These are mostly based on interior point methods and are relatively efficient, though they cannot perform deformation with many degrees of freedom in real time rates and none of these

is directly suitable for a GPU implementation. Our method outperforms these by a large margin. For example, in Figure 2, we show a speedup of $\times 125$ when comparing our GPU Newton solver with the Mosek interior point solver that was used in [Chen and Weber 2015]. To the best of our knowledge, the only bounded distortion deformation method that utilizes the GPU is [Hefetz et al. 2017] which projects maps to the bounded distortion space by using alternating tangential projections (ATP). It is extremely fast and provably converges to a bounded isometric-and-conformal distortion map with any user-specified bounds, albeit ATP does not handle positional constraints, nor minimize user-defined energies.

The Newton's algorithm (see [Nocedal and Wright 2006] for a comprehensive review) is an established method for the minimization of smooth nonlinear objective functions, and is widely known for its good convergence properties due to its second order nature. Each Newton iteration is very effective in reducing the energy, and typically a small amount of iterations is needed to reach low energy. For the Newton iteration to produce a descent direction, the Hessian of the energy must be positive definite and since the energies of interest are nonconvex in general, several strategies for modifying the Hessian to become positive definite have been developed (see [Nocedal and Wright 2006, Chapter 3.4] for a concise overview).

The most straightforward approach is to compute the eigenfactorization (also known as spectral decomposition) and to replace all the negative eigenvalues with sufficiently small positive values. The obtained modified Hessian is the closest (in Frobenius norm) positive definite matrix to the unmodified one. We compare our results with this method and show our superiority. Another common approach [Fu and Liu 2016] is to shift the spectrum by trying iteratively to add a positive constant to the diagonal of the Hessian until the Cholesky factorization succeeds (which is an indication for positive definiteness). In practice though, such modifications are expensive to compute.

A more affordable approach is to decompose the Hessian as a sum of individual per-element matrices and to perform the modification on each element separately. The approximated Hessian is then obtained by a summation of the modified per-element matrices. In [Teran et al. 2005], each per-element tensor is encoded as a 9×9 block diagonal matrix and in [Fu and Liu 2016], 4×4 or 9×9 matrices are constructed and modified numerically. This approach is in general more efficient than a full Hessian modification but is still time consuming due to the large number of required matrix factorizations. Our approach is conceptually similar, however, due to the infinite support of the harmonic basis functions, our per-element Hessians are dense $4n \times 4n$ matrices (n is typically of order of hundreds). Performing the factorization numerically is practically impossible. A key contribution (Section 7) in our paper is that each such $4n \times 4n$ factorization can be computed analytically. Furthermore, we show that the number of elements (samples) needed in practice to achieve good convergence rates can be relatively small, ensuring that the assembly of the individual Hessian contributions does not dominate the solver performance. The computation of the Hessian modification (as well as all the other parts of our algorithm) is done in parallel and is perfectly suitable for implementation on graphics hardware (Appendix G) leading to an algorithm that is orders of magnitude faster than state-of-the-art methods.

If the objective function is a nonlinear sum of squares, the Gauss-Newton (GN) or the Levenberg-Marquardt (LM) algorithms can be used. These use first order derivatives only, sidestepping the requirement to compute-and-modify the Hessian matrix. [Vaxman et al. 2015] uses LM to compute conformal deformations in 2-and-3 dimensions. Another class of first order methods which can be applied to more general type of objectives are Quasi-Newton algorithms. BFGS for example, incrementally approximate the Hessian based on the gradients and the updates from previous iterations. L-BFGS further avoids inverting the Hessian by directly approximating its inverse from k recent iterations [Liu et al. 2017]. A single iteration of L-BFGS is then more efficient than that of a second order method, but the number of iterations is typically much larger. [Smith and Schaefer 2015] used L-BFGS to optimize a particular choice of isometric energy. We compare our solver for this particular energy against L-BFGS, showing a dramatic reduction in the number of iterations and increase in speed.

The Accelerated Quadratic Proxy (AQP) method is another notable first order method [Kovalsky et al. 2016] that can optimize isometric-like energies. A main observation made by Kovalsky et al. is that the Hessian of the particular deformation energy can be approximated by the mesh Laplacian matrix which is the Hessian of the quadratic Dirichlet energy. In contrast to a Newton iteration which requires a full linear solve at each iteration, AQP merely requires a backward and forward substitutions against the factor of the *fixed* Laplacian matrix, leading to extremely fast iterations. In order to reduce the number of iterations, which is typically much larger than that of Newton's, an acceleration technique was used. Albeit, this advocates the use of hard positional constraints. Since finding a locally injective map satisfying hard positional constraints is a notoriously hard problem, AQP only produces orientation preserving maps which is a weaker requirement. To obtain an initialization, AQP first computes an arbitrary map satisfying the positional constraints, and then projects it using the method of [Kovalsky et al. 2015] to the space of orientation preserving maps. While [Kovalsky et al. 2015] is a relatively efficient and robust method, it does not guarantee convergence and utilizing it as initialization extends the overall computation time. The method of [Martin et al. 2013] takes a signal processing viewpoint for mesh optimization where powers of the Laplacian matrix are utilized.

Recently, Rabinovich et al. [2017] proposed Scalable Locally Injective Mappings (SLIM) for optimizing rotation-invariant mesh-based energies via a modification of the local-global ARAP algorithm [Liu et al. 2008]. Unlike AQP which uses a fixed Hessian approximation, SLIM reweights a proxy energy iteratively in a way that aligns its gradient with that of the true energy. While AQP and SLIM are more general and are not restricted to planar maps, we demonstrate that our solver is significantly faster, with much shorter per-iteration time and dramatically smaller number of iterations. In contrast to all the above mentioned mesh-based methods, which are based on sparse linear algebra, our method operates within a reduced dimensional subspace which is formulated in dense linear algebra terms. While the sparse direct linear solvers utilized in these mesh-based algorithms exploit multi-core CPU parallelism, existing implementations on graphics hardware (e.g. SuiteSparse, cuSolverSP) are still immature.

3 BACKGROUND

For completeness we include a short introduction to some basic concepts in optimization, complex analysis, and planar harmonic maps. For further reading we refer to the comprehensive book by Nocedal and Wright [2006] for optimization related topics, the book by Ahlfors for complex analysis related topics [1979], and the concise book by Duren regarding harmonic planar maps [2004].

3.1 Newton's Method

Given a scalar function $E(\mathbf{x})$ of a vector variable $\mathbf{x} \in \mathbb{R}^n$, Newton's method searches for a stationary point of $E(\mathbf{x})$ by iteratively approximating it using its second order Taylor expansion:

$$E(\mathbf{x}) = E(\mathbf{x}_k + \Delta\mathbf{x}) \approx E(\mathbf{x}_k) + \Delta\mathbf{x}^\top \nabla E(\mathbf{x}_k) + \frac{1}{2} \Delta\mathbf{x}^\top \nabla^2 E(\mathbf{x}_k) \Delta\mathbf{x},$$

where $\Delta\mathbf{x} = \mathbf{x} - \mathbf{x}_k \in \mathbb{R}^n$ is the search direction, and $\nabla E \in \mathbb{R}^n$ and $\nabla^2 E \in \mathbb{R}^{n \times n}$ denotes the gradient and the Hessian of E respectively. This second order polynomial in $\Delta\mathbf{x}$ has vanishing derivative when the following linear system is satisfied:

$$\nabla^2 E(\mathbf{x}_k) \Delta\mathbf{x} = -\nabla E(\mathbf{x}_k). \quad (1)$$

The direction $\Delta\mathbf{x}$ decreases the energy only if the second order polynomial is convex, or equivalently, the Hessian matrix $\nabla^2 E$ is positive definite. Hence, when needed, the Hessian is modified. The solution to the linear system produces the update for the current iteration $\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta\mathbf{x}$. To ensure convergence (to a local minimum) from an arbitrary starting point, a line search along the search direction $\Delta\mathbf{x}$ should be used to find a step size t , that satisfies the Wolfe conditions [Nocedal and Wright 2006, Chapter 3.1].

3.2 Complex-valued Functions as Planar Maps

A real-valued function $u(x, y) : \Omega \rightarrow \mathbb{R}$, is harmonic if it satisfies the Laplace equation $u_{xx} + u_{yy} = 0$. A planar harmonic map is a map $f : \Omega \rightarrow \mathbb{R}^2$ where $f = [u(x, y), v(x, y)]$, and the component functions u and v are harmonic. When we identify \mathbb{R}^2 with the complex plane \mathbb{C} using complex notation $z = x + iy$, we can express f as a complex-valued function $f(z) = u(z) + iv(z)$. Thus a complex-valued harmonic function $f : \Omega \subset \mathbb{C} \rightarrow \mathbb{C}$ can be interpreted as a harmonic planar map. A complex-valued function $f = u + iv$ is said to be *holomorphic* if it satisfies the Cauchy-Riemann equations. $f(z)$ is called *anti-holomorphic* if $\bar{f}(\bar{z})$ is holomorphic.

The Wirtinger derivatives of a complex-valued function are defined as follows: $f_z := \frac{1}{2}(f_x - if_y)$ and $f_{\bar{z}} := \frac{1}{2}(f_x + if_y)$. Using the Wirtinger derivatives, the Cauchy-Riemann equations are expressed concisely as $f_{\bar{z}} \equiv 0$. The Laplacian $f_{xx} + f_{yy}$ is concisely expressed as $4f_{z\bar{z}}$, hence the condition for harmonicity (of $\text{Re} f$ and $\text{Im} f$) is simply $f_{z\bar{z}} \equiv 0$. It is then easy to see that the real and imaginary components of a holomorphic function are harmonic. The same is true for anti-holomorphic functions. The converse is not true since $f_{z\bar{z}} \equiv 0$ does not imply $f_z \equiv 0$ nor $f_{\bar{z}} \equiv 0$.

The Wirtinger derivatives are useful for defining many of the local geometric distortion quantities of a C^1 planar map. We have that the local change in area is: $\det(J_f) = |f_z|^2 - |f_{\bar{z}}|^2$, where J_f is the 2×2 Jacobian matrix of the map. A C^1 map f is locally injective and orientation preserving at a neighborhood of a point z if $\det(J_f(z)) > 0$, or alternatively if $|f_z(z)| > |f_{\bar{z}}(z)|$.

Isometric and conformal distortion measures are typically formulated in terms of the singular values of J_f , $0 \leq \sigma_2 \leq \sigma_1$:

$$\sigma_1(z) = |f_z| + |f_{\bar{z}}|, \quad \sigma_2(z) = \left| |f_z| - |f_{\bar{z}}| \right|. \quad (2)$$

Note that for orientation preserving locally injective maps, the expression on the right simplifies to $\sigma_2(z) = |f_z| - |f_{\bar{z}}| > 0$. The little dilatation $k(z) = |f_{\bar{z}}(z)| / |f_z(z)|$ is a measure of conformal distortion. We discuss isometric distortion measures in Section 6.

4 HARMONIC MAPS ON MULTIPLY-CONNECTED DOMAINS

As mentioned in the Introduction, our work generalizes existing techniques that are designed for simply-connected domains to multiply-connected domains. We base our construction on the following theorem which we prove in Appendix A.

THEOREM 4.1 (HARMONIC DECOMPOSITION). *Let Ω be a multiply-connected planar domain with N holes K_1, \dots, K_N , and choose N arbitrary points ρ_i in K_i . Then, any harmonic map $f : \Omega \rightarrow \mathbb{C}$ can be represented as:*

$$f(z) = \tilde{\Phi}(z) + \overline{\tilde{\Psi}(z)} + \sum_{i=1}^N \omega_i \ln |z - \rho_i|, \quad (3)$$

where $\tilde{\Phi}, \tilde{\Psi} : \Omega \rightarrow \mathbb{C}$ are holomorphic functions, and $\omega_1, \dots, \omega_N$ are some complex coefficients.

The harmonic decomposition (3) is unique up to an additive complex constant that can be chosen by setting e.g. $\tilde{\Psi}(z_0) = 0$ for an arbitrary point $z_0 \in \Omega$. Theorem 4.1 extends a similar result for simply-connected domains for which the summation term in (3) is missing (c.f. [Duren 2004] Section 1.2). While $\tilde{\Phi}(z) + \tilde{\Psi}(z)$ is still harmonic on multiply-connected domains, the additional summation term is essential to represent certain harmonic maps and without it, the representation is incomplete.

4.1 Locally Injective Harmonic Maps

The main theoretical result in [Chen and Weber 2015] is a boundary-value characterization of the injectivity and the conformal-and-isometric distortion of planar harmonic maps on simply-connected domains. We generalize this result to the multiply-connected case (see Appendix B for the proof):

THEOREM 4.2 (BOUNDED DISTORTION). *A planar harmonic map $f : \Omega \rightarrow \mathbb{C}$ on a multiply-connected domain with exterior boundary curve γ_0 oriented counterclockwise and interior boundary curves $\gamma_1 \cdots \gamma_N$ oriented clockwise, is locally injective with an upper bound $k \in [0, 1)$ on the conformal distortion, a lower bound $\sigma_2 > 0$ on the small singular value of the Jacobian, and an upper bound $\sigma_1 < \infty$ on the large singular value at every point z in Ω if and only if:*

$$\oint_{\gamma_0} \frac{f'_z(w)}{f_z(w)} dw + \sum_{i=1}^N \oint_{\gamma_i} \frac{f'_z(w)}{f_z(w)} dw = 0, \quad (4a)$$

$$0 \leq k(w) \leq k \quad \forall w \in \partial\Omega, \quad (4b)$$

$$\sigma_1(w) \leq \sigma_1 \quad \forall w \in \partial\Omega, \quad (4c)$$

$$\sigma_2 \leq \sigma_2(w) \quad \forall w \in \partial\Omega. \quad (4d)$$




Fig. 3. A knot. Left: doubly-connected domain Ω . Right: an image of a map $f : \Omega \rightarrow \mathbb{C}$. While f is locally injective, f and the identity are not regular homotopic. Intuitively, one cannot “morph” the annulus on the left to the one on the right continuously without breaking local injectivity. The conditions of Theorem 4.2 on f hold. The integral (Equation (4a)) over the exterior boundary is $2\pi i$ while the integral over the interior one is $-2\pi i$.

Intuitively, Theorem 4.2 states that in order to induce global distortion bounds, it is sufficient to bound the distortion of a harmonic map on the *boundary* curves γ_i as long as (4a) holds. The boundary integral condition (4a) is equivalent to the condition that f_z is non-vanishing throughout the domain. Namely that $f_z(z) \neq 0$ (Appendix B). In contrast to the simply-connected domain, the integral over the exterior boundary γ_0 may be nonzero. This gives rise to maps which are not regular homotopic to the identity map but can still be locally injective (see Figure 3).

As opposed to [Chen and Weber 2015], our deformation algorithm does not *explicitly* impose user-defined bounds k, σ_1, σ_2 on the distortion. Instead, we favor an overall lower *average* distortion. By using energies that become infinite when the map degenerates at a boundary point, a finite bound on distortion is naturally formed on the boundary. Theorem 4.2 then assures that this bound is also global. At each iteration, our algorithm certifies the map as locally injective by verifying that (4a) holds, as well as the following:

$$|f_z(w)| > |f_{\bar{z}}(w)| \quad \forall w \in \partial\Omega. \quad (5)$$

In Section 8 we explain how to enforce (4a) and (5) in practice.

5 DISCRETIZATION

We discretize Equation (3) to obtain a finite dimensional space of harmonic maps. The holomorphic functions $\tilde{\Phi}$ and $\tilde{\Psi}$ are discretized by using the Cauchy complex barycentric coordinates [Weber et al. 2009] which are derived by approximating the boundary of the domain using a polygonal shape. Let $\hat{P} = \{z_1, z_2, \dots, z_m\}$, $z_j \in \mathbb{C}$ be the vertices of a multiply-connected planar polygon (the cage). The vertices of the cage are split into sets such that the first set represents the exterior polygon oriented counterclockwise while the following sets represent the interior polygons oriented clockwise (see Figure 4).




Fig. 4. Notations. The purple area is a doubly-connected domain Ω bounded by a polygon P with 11 vertices. The solid black polygon is an outward offset cage $\hat{P} = \gamma_0 \cup \gamma_1$ with vertices z_j (cyan) composed of exterior polygon γ_0 oriented counterclockwise, and interior clockwise polygon γ_1 . The point z (green) is in Ω or on $\partial\Omega$. $A_j = z_j - z_{j-1}$, $B_j(z) = z_j - z$, and $[v_i, v_{i+1}]$ is a small segment (orange) on P . The point ρ_1 (red) is in γ_1 .

The holomorphic functions $\tilde{\Phi}$ and $\tilde{\Psi}$ are represented as:

$$\tilde{\Phi}(z) = \sum_{j=1}^m \tilde{C}_j(z) \varphi_j, \quad \tilde{\Psi}(z) = \sum_{j=1}^m \tilde{C}_j(z) \psi_j, \quad (6)$$

where $\tilde{C}_j(z)$ is the j^{th} holomorphic Cauchy barycentric coordinate associated with vertex z_j and φ_j, ψ_j are complex coefficients. $\tilde{C}_j(z)$ possess a rather simple closed-form expression (see Appendix D). Note that while Weber et al. [2009] assumed that \hat{P} is simply-connected, their derivation can be extended to the case of multiple boundary components due to the fact that Cauchy’s integral formula still holds. This can be done by integrating the Cauchy kernel over *all* the boundary components. Practically, representing holomorphic functions on such a domain simply boils down to using additional basis functions that correspond to vertices that lie on the interior boundaries. We note, that the complex barycentric map induced by the Cauchy coordinates on multiply-connected domains has a non-empty null space of (complex) dimension $2N$ which corresponds to a similarity transformation of the vertices of the interior boundary polygons. To remove this ambiguity, we fix the first two complex coefficients of each hole to zero.

We proceed by substituting the holomorphic functions $\tilde{\Phi}$ and $\tilde{\Psi}$ in the harmonic decomposition (3) with the expressions from (6). The additional term $\sum_{i=1}^N \omega_i \ln |z - \rho_i|$ in (3) is substituted with the expression:

$$\sum_{j=m+1}^{m+N} (\varphi_j + \overline{\psi_j}) \ln |z - \rho_{j-m}|, \quad (7)$$

where φ_j, ψ_j are $2N$ additional complex coefficients. Note that splitting ω_i into two variables creates some unnecessary redundancy, nevertheless it greatly simplifies the exposition as well as the implementation. Later on, we will remove this redundancy by enforcing a constraint $\varphi_j = \overline{\psi_j}$, $j = m+1, \dots, m+N$ for each hole. Finally, rearranging terms and denoting $n = m+N$ leads to:

$$f(z) = \sum_{j=1}^n C_j(z) \varphi_j + \sum_{j=1}^n \overline{C_j(z)} \psi_j, \quad (8)$$

$$C_j(z) = \begin{cases} \tilde{C}_j(z) & j = 1, \dots, m \\ \ln |z - \rho_{j-m}| & j = m+1, \dots, n. \end{cases}$$

Equation (8) is our closed-form expression for a finite dimensional harmonic subspace with $2n$ complex variables $\{\varphi_j, \psi_j\}_{j=1}^n$. The dimension of the (complex) null space is $4N + N + 1$, where $4N$ is due to the 2 complex DOF per hole for each of the two Cauchy barycentric maps (as explained above), N is due to the redundancy introduced by Equation (7) and the term 1 is due to the constant DOF of the harmonic decomposition (3).

We proceed by constructing simple formulas for the Wirtinger derivatives of our harmonic map. Luckily, just like the map itself, it can be expressed in closed-form. By linearity of the Wirtinger operators, the fact that the derivative of the Cauchy coordinates with respect to \bar{z} is 0, and assuming that $\varphi_j = \bar{\psi}_j, \forall j = m + 1, \dots, n$ we have:

$$f_z(z) = \sum_{j=1}^n D_j(z)\varphi_j, \quad \bar{f}_z(z) = \sum_{j=1}^n D_j(z)\psi_j, \quad (9)$$

$$D_j(z) = \begin{cases} \tilde{D}_j(z) & j = 1, \dots, m \\ \frac{1}{z - \rho_{j-m}} & j = m + 1, \dots, n. \end{cases}$$

The derivatives of the Cauchy barycentric coordinates $\tilde{D}_j(z)$ (Appendix D) are holomorphic and so is $\frac{1}{z - \rho_{j-m}}$, hence it is clear that f_z and \bar{f}_z are both holomorphic (though not necessarily integrable).

6 ISOMETRIC ENERGIES

The main approach in interactive shape deformation is to design and minimize a distortion energy that aggregates a differential quantity that strives to keep the map as-isometric-as-possible. Such differential quantities are minimized when the map is locally isometric. However, since a perfect isometry cannot be obtained in general in the presence of positional constraints, the particular definition of proximity to isometry greatly affects the end result. It is a common practice to measure distortion at a point z as a function $E(z) = E(\sigma_1(z), \sigma_2(z))$ of the two singular values of the Jacobian matrix $J_f(z)$. Such a function is invariant to rigid motions of both the source and the target domains [Rabinovich et al. 2017]. $E(z)$ should be minimized (at a single point z) if $\sigma_1(z) = \sigma_2(z) = 1$. Energies that minimize *conformal* distortion are also popular. Our harmonic subspace contains many pure (with zero distortion) conformal maps, hence, if desired, we can simply restrict the subspace by eliminating the ψ_j variables and use any of our isometric energies to regularize the result (Figure 5).

A popular choice in graphics is the as-rigid-as-possible (ARAP) energy [Igarashi et al. 2005; Liu et al. 2008; Sorkine and Alexa 2007] which defines local proximity to isometry via $\|J - R\|_F^2$, where $\|\cdot\|_F$ denotes the Frobenius norm, and R denotes the closest rotation to J . It can be expressed as $E_{\text{ARAP}} = (\sigma_1 - 1)^2 + (\sigma_2 - 1)^2$. The main limitation of ARAP is that it favors shrinkage over expansion, and in particular it stays finite even if the map is degenerated ($\sigma_2(z) = 0$). Hence, ARAP tends to attract the map to a configuration which is not locally injective. Our solver and derivation are applicable to E_{ARAP} , albeit ARAP is arguably a poor choice of energy if local injectivity is sought after.

Appropriate choices of isometric energy for designing locally injective maps become infinite when the map collapses locally, which serves as a natural barrier term. Since our optimization depends on




Fig. 5. Harmonic deformations of a triply-connected domain. The domain on the left has $N = 2$ holes in it. The result in the middle minimizes the Advanced MIPS energy, while the one on the right is a pure conformal map with least isometric distortion.

high order derivatives, we focus on energies which are smooth in our variables φ_j, ψ_j . To this end, we express the distortion measure at z as a smooth function of $|f_z|^2$ and $|\bar{f}_z|^2$. Since $|f_z|^2$ and $|\bar{f}_z|^2$ are quadratic functions in φ_j, ψ_j , the composition is also smooth. The symmetric Dirichlet isometric energy $E_{\text{iso}} = \frac{1}{2}|J|_F^2 + \frac{1}{2}|J|_F^{-2}$ has been successfully used in [Kovalsky et al. 2016; Rabinovich et al. 2017; Schreiner et al. 2004; Smith and Schaefer 2015] for mesh parameterization. It can be expressed in terms of the singular values:

$$E_{\text{iso}}(\sigma_1, \sigma_2) = \frac{1}{2} (\sigma_1^2 + \sigma_1^{-2} + \sigma_2^2 + \sigma_2^{-2}) \quad (10a)$$

$$= \frac{\sigma_1^2 + \sigma_2^2}{2} \left(1 + \frac{1}{\sigma_1^2 \sigma_2^2} \right). \quad (10b)$$

Recall that for an orientation preserving locally injective planar map f , we have $\sigma_1 = |f_z| + |\bar{f}_z|$ and $\sigma_2 = |f_z| - |\bar{f}_z|$. We get:

$$\frac{\sigma_1^2 + \sigma_2^2}{2} = |f_z|^2 + |\bar{f}_z|^2, \quad \sigma_1 \sigma_2 = |f_z|^2 - |\bar{f}_z|^2, \quad (11)$$

where the left term is the Dirichlet energy and the right term is $\det(J)$. Substituting into Equation (10b) gives:

$$E_{\text{iso}}(|f_z|^2, |\bar{f}_z|^2) = (|f_z|^2 + |\bar{f}_z|^2) \left(1 + \frac{1}{(|f_z|^2 - |\bar{f}_z|^2)^2} \right),$$

which is clearly smooth in $|f_z|^2$ and $|\bar{f}_z|^2$ since $|f_z|^2 \neq |\bar{f}_z|^2$.

In Table 1, we list several possible choices for such smooth isometric energies. These include the exponential symmetric Dirichlet isometric energy: $E_{\text{exp}} = \exp(s \cdot E_{\text{iso}})$ [Rabinovich et al. 2017]. The motivation to use E_{exp} is to penalize more drastically high values of the distortion measure. As demonstrated in Figure 6, this allows the user to trade-off low average versus low maximal distortion. Finally, we include the Advanced MIPS energy [Fu et al. 2015] which provides user-controlled balance between area preservation and angle preservation.

We define the isometric distortion of a planar harmonic map f as a boundary integral over the pointwise distortion quantity $E(w)$:

$$E^f = \oint_{\partial\Omega} E(w) ds. \quad (12)$$

Table 1. Smooth isometric energies. First column: isometric measure in terms of the singular values. Second column: isometric measure in terms of $|f_z|^2$, $|f_{\bar{z}}|^2$ which are denoted x , y for clarity. It is straightforward to see that these expressions are smooth in x , y wherever $x \neq y$, which is always the case for our locally injective maps. Third column: the parameters that are needed for instantiating the gradient (17) and the Hessian (18) expressions. Forth column: The four eigenvalues of \mathbf{K} (not necessarily sorted). Potentially negative ones are highlighted.

	$E(\sigma_1, \sigma_2)$	$E(x, y) = E(f_z ^2, f_{\bar{z}} ^2)$	$(\alpha_1, \alpha_2, \beta_1, \beta_2, \beta_3)^T$	Eigenvalues
E_{iso}	$\frac{1}{2} (\sigma_1^2 + \sigma_2^2 + \sigma_1^{-2} + \sigma_2^{-2})$	$E_{\text{iso}} = (x + y) \left(1 + \frac{1}{(x-y)^2} \right)$	$\begin{pmatrix} \alpha_1^{\text{iso}} = 1 - (x + 3y)(x - y)^{-3} \\ \alpha_2^{\text{iso}} = 1 + (3x + y)(x - y)^{-3} \\ \beta_1^{\text{iso}} = 2(x + 5y)(x - y)^{-4} \\ \beta_2^{\text{iso}} = 2(5x + y)(x - y)^{-4} \\ \beta_3^{\text{iso}} = -6(x + y)(x - y)^{-4} \end{pmatrix}$	$\underline{2\alpha_1}, 2\alpha_2, \lambda_3, \lambda_4$
E_{exp}	$\exp\left(\frac{s}{2} (\sigma_1^2 + \sigma_2^2 + \sigma_1^{-2} + \sigma_2^{-2})\right)$	$\exp\left(s(x + y) \left(1 + \frac{1}{(x-y)^2} \right)\right)$	$s \text{Exp} \begin{pmatrix} \alpha_1^{\text{iso}} \\ \alpha_2^{\text{iso}} \\ \beta_1^{\text{iso}} + s(\alpha_1^{\text{iso}})^2 \\ \beta_2^{\text{iso}} + s(\alpha_2^{\text{iso}})^2 \\ \beta_3^{\text{iso}} + s\alpha_1^{\text{iso}}\alpha_2^{\text{iso}} \end{pmatrix}$	$\underline{2\alpha_1}, 2\alpha_2, \lambda_3, \lambda_4$
E_{AMIPS}	$\exp\left(s \left(\frac{\sigma_1}{\sigma_2} + \frac{\sigma_2}{\sigma_1} \right) + (\sigma_1\sigma_2 + \sigma_1^{-1}\sigma_2^{-1})\right)$	$\exp\left(\frac{2s(x+y)+1}{x-y} + x - y\right)$	$\begin{pmatrix} (1 - (4sy + 1)(x - y)^{-2}) E_{\text{AMIPS}} \\ (-1 + (4sx + 1)(x - y)^{-2}) E_{\text{AMIPS}} \\ \alpha_1^2 E_{\text{AMIPS}}^{-1} + 2(4sy + 1)(x - y)^{-3} E_{\text{AMIPS}} \\ \alpha_2^2 E_{\text{AMIPS}}^{-1} + 2(4sx + 1)(x - y)^{-3} E_{\text{AMIPS}} \\ \alpha_1 \alpha_2 E_{\text{AMIPS}}^{-1} - (4sx + 4sy + 2)(x - y)^{-3} E_{\text{AMIPS}} \end{pmatrix}$	$\underline{2\alpha_1}, \underline{2\alpha_2}, \underline{\lambda_3}, \underline{\lambda_4}$

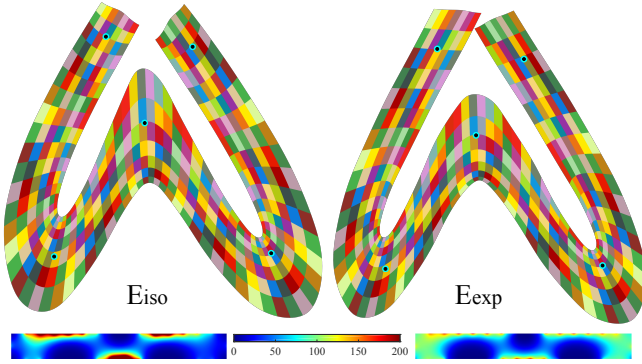


Fig. 6. Comparison of E_{iso} (left) and E_{exp} (right). A straight bar shape is deformed into an extreme pose. Exponentiating the E_{iso} distortion measure reduces the maximal distortion (red color) but increases the average distortion. Both color maps show E_{iso} .

We use the superscript f throughout the paper to denote the boundary integrated distortion and omit it to denote the pointwise quantity. We also experimented with an area integral instead of the boundary one, but did not notice any benefit. Theorem 4.2 ensures that for any bound on $E(w)$ along the boundary, a global upper bound on $\sigma_1(z)$, and a global lower bound on $\sigma_2(z)$ exist. Hence, a global bound on $E(z)$ is naturally formed.

The gradient and the Hessian of the overall isometric energy are:

$$\nabla E^f = \oint_{\partial\Omega} \nabla E(w) ds \quad (13)$$

$$\nabla^2 E^f = \oint_{\partial\Omega} \nabla^2 E(w) ds. \quad (14)$$

The gradient and the Hessian of our energy measures have relatively simple closed-form expressions. A complete derivation is given in Appendix E. Let $D = (D_1, D_2, \dots, D_n) \in \mathbb{C}^{1 \times n}$ be a complex row vector, where D_j is defined as in (9). We use bold symbols to denote

real vectors and matrices. Define the real matrix D (note the bold symbol) as:

$$\mathbf{D} = \begin{bmatrix} \text{Re}(D) & -\text{Im}(D) \\ \text{Im}(D) & \text{Re}(D) \end{bmatrix} \in \mathbb{R}^{2 \times 2n}. \quad (15)$$

We express the complex Wirtinger derivatives as 2×1 real vectors:

$$f_z = \begin{bmatrix} \text{Re}(f_z) \\ \text{Im}(f_z) \end{bmatrix}, \quad \overline{f_z} = \begin{bmatrix} \text{Re}(\overline{f_z}) \\ \text{Im}(\overline{f_z}) \end{bmatrix} \in \mathbb{R}^{2 \times 1}. \quad (16)$$

The gradient of E with respect to the $4n$ real variables is:

$$\nabla E(z) = 2 \begin{bmatrix} \alpha_1 \mathbf{D}^T f_z \\ \alpha_2 \mathbf{D}^T \overline{f_z} \end{bmatrix} \in \mathbb{R}^{4n \times 1}, \quad (17)$$

where α_1, α_2 are real parameters which depend on the particular choice of energy. Table 1 provides the parameters needed to instantiate some particular energies (out of many possible). The expression for the $4n \times 4n$ Hessian of $E(z)$ at a single point is:

$$\nabla^2 E(z) = \underbrace{\begin{bmatrix} \mathbf{D}^T & 0 \\ 0 & \mathbf{D}^T \end{bmatrix}}_{4n \times 4} \underbrace{\mathbf{K}}_{4 \times 4} \underbrace{\begin{bmatrix} \mathbf{D} & 0 \\ 0 & \mathbf{D} \end{bmatrix}}_{4 \times 4n} \in \mathbb{R}^{4n \times 4n}, \quad (18)$$

where \mathbf{K} is the a 4×4 real matrix:

$$\mathbf{K} = \begin{bmatrix} 2\alpha_1 I + 4\beta_1 f_z f_z^T & 4\beta_3 f_z \overline{f_z}^T \\ 4\beta_3 \overline{f_z} f_z^T & 2\alpha_2 I + 4\beta_2 \overline{f_z} \overline{f_z}^T \end{bmatrix} \in \mathbb{R}^{4 \times 4}. \quad (19)$$

The parameters $\beta_1, \beta_2, \beta_3$ are also energy dependent (see Table 1).

6.1 Positional Constraints

The point-to-point (P2P) metaphor is an intuitive drag-and-drop user-interface that best fits interactive deformation tasks. The user of our system can add or remove (by clicking) positional constraints at any time during interaction. Dragging the P2P handles signals the application to invoke the optimization and to render the updated result. We incorporate the P2P constraints as soft constraints as was

advocated by [Chen and Weber 2015; Poranne and Lipman 2014; Rabinovich et al. 2017]. We define the P2P energy:

$$E_{\text{p2p}}^f = \frac{1}{2} \sum_{i=1}^{|\mathcal{P}|} |f(p_i) - q_i|^2,$$

where $p_i \in \mathcal{P} \subset \Omega$ is the position of the i^{th} handle in the source domain, and $q_i \in \mathbb{C}$ is its target desired position. The gradient and the Hessian of E_{p2p}^f are derived in Appendix F. The full energy of our unconstrained minimization problem is:

$$E_{\text{Def}}^f = E^f + \lambda E_{\text{p2p}}^f, \quad (20)$$

where λ is a user-defined weight that balances the two terms.

7 POSITIVE DEFINITE HESSIAN

The Hessian of E_{p2p}^f is positive semi-definite (PSD). However, the Hessian of E^f is, in general, not PSD, and neither is the Hessian of E_{Def}^f (Equation (20)). Our key observation here is that the closest (in Frobenius norm) PSD matrix to the Hessian $\nabla^2 E(z) \in \mathbb{R}^{4n \times 4n}$ at a *single* point z , can be expressed in closed-form. With that at hand we substitute Equation (14) with:

$$\nabla^2 E^{f+} = \oint_{\partial\Omega} \nabla^2 E^+(w) ds, \quad (21)$$

where $\nabla^2 E^+(w)$ is the closest PSD matrix to $\nabla^2 E(w)$. Since, the integral (or sum) of PSD matrices is PSD (due to the convex cone structure of the PSD set), it is clear that the modified Hessian of the isometric energy (21) is guaranteed to be PSD. In Appendix I, we further show that the dimension of the null space of (21) in the case of E_{iso} and E_{exp} is 2 or 3. Therefore with 2 or more positional constraints, $\nabla^2 E^{f+}$ is nonsingular. We refer to the variant of Newton's method that computes the closest PSD matrix to $\nabla^2 E^f$ as Newton-Eigen and stress that our modified Newton iterations are dramatically faster to compute. Moreover, throughout extensive experiments, we observe that our local modification leads to iterations which are more *effective*, where less iterations are required for convergence (Figures 1, 10, 11).

7.1 Hessian Modification

Our goal is to compute the eigen-factorization of the local Hessian matrix $\nabla^2 E$. On the one hand, doing it numerically is impractical, whereas on the other hand, analytic eigen-factorization of a $4n \times 4n$ matrix is challenging to derive. Our first observation is that each of the nontrivial eigenvectors (with nonzero eigenvalue) of $\nabla^2 E$ can be expressed as a product of $\mathbf{B} = \begin{bmatrix} D & 0 \\ 0 & D \end{bmatrix}$ and a corresponding eigenvector of \mathbf{K} (Equation (19)). This is due to the fact that \mathbf{B} has 4 rows that are orthogonal to each other (easy to verify by computing the inner product of each pair) and all the rows have the same norm. Furthermore, the eigenvalues of \mathbf{K} and $\nabla^2 E$ are the same up to a positive scale. Hence, we have reduced our problem to that of analytically computing the eigenvalues of a 4×4 matrix. This is still challenging as these are the roots of a 4^{th} order polynomial. Denote the (unsorted) eigenvalues of \mathbf{K} as $(\lambda_1, \lambda_2, \lambda_3, \lambda_4)$. Our second observation is that $\lambda_1 = 2\alpha_1$ and $\lambda_2 = 2\alpha_2$. To see why $2\alpha_1$ is an

eigenvalue, subtract $2\alpha_1$ from the diagonal of \mathbf{K} . The first two rows of $\mathbf{K} - 2\alpha_1 I$ are:

$$\underbrace{f_z}_{2 \times 1} \underbrace{\begin{bmatrix} 4\beta_1 f_z^T & 4\beta_3 \overline{f_z}^T \end{bmatrix}}_{1 \times 4}, \quad (22)$$

where we can see that the 1×4 row vector on the right hand side of (22) is multiplied by two scalars (the elements of f_z), hence, these two rows are linearly dependent, meaning that the matrix $\mathbf{K} - 2\alpha_1 I$ is singular and $2\alpha_1$ is a root of the characteristic polynomial. Showing that $2\alpha_2$ is an eigenvalue can be done similarly. Knowing that $2\alpha_1$ and $2\alpha_2$ are two eigenvalues of \mathbf{K} , we can compute the other two eigenvalues by directly solving the quartic characteristic equation. The derivation is long but straightforward, hence omitted. We get these expressions:

$$\lambda_{3,4} = s_1 \pm \sqrt{s_2^2 + 16\beta_3^2 |f_z|^2 |f_z|^2}, \quad (23)$$

where $s_{1,2} = \alpha_1 + 2\beta_1 |f_z|^2 \pm (\alpha_2 + 2\beta_2 |f_z|^2)$.

The corresponding four eigenvectors are:

$$\begin{pmatrix} \text{Im}(f_z), & -\text{Re}(f_z), & 0, & 0 \\ 0, & 0, & \text{Im}(\overline{f_z}), & -\text{Re}(\overline{f_z}) \\ \text{Re}(f_z), & \text{Im}(f_z), & t_1 \text{Re}(\overline{f_z}), & t_1 \text{Im}(\overline{f_z}) \\ \text{Re}(f_z), & \text{Im}(f_z), & t_2 \text{Re}(\overline{f_z}), & t_2 \text{Im}(\overline{f_z}) \end{pmatrix} \quad (24)$$

where:

$$t_{1,2} = \frac{\lambda_{3,4} - 2\alpha_1 - 4\beta_1 |f_z|^2}{4\beta_3 |f_z|^2}.$$

The signs of these eigenvalues depend on the particular choice of isometric energy, therefore in the most general case, the 4 eigenvalues are evaluated and negative ones are substituted with 0 to obtain the modified Hessian. For particular energy choices, it is possible to simplify matter even more. For the first two energies listed in Table 1, we have that only $\lambda_1 = 2\alpha_1$ can be (at times) negative. This allows us to directly express the modified matrix. For example, for E_{iso} , it turns out that λ_3, λ_4 have quite simple expressions:

$$\lambda_3 = 4(1 + 3(|f_z| + |\overline{f_z}|)^{-4}), \quad \lambda_4 = 4(1 + 3(|f_z| - |\overline{f_z}|)^{-4}).$$

The spectrum is then sorted as follows: $2\alpha_1 \leq \lambda_3 \leq 2\alpha_2 \leq \lambda_4$. With this information at hand, the modification is done by checking for the sign of α_1 , and if it is negative, we substitute \mathbf{K} in (18) with:

$$\mathbf{K}^+ = \begin{bmatrix} (\frac{2\alpha_1}{|f_z|^2} + 4\beta_1) f_z f_z^T & 4\beta_3 f_z \overline{f_z}^T \\ 4\beta_3 \overline{f_z} f_z^T & 2\alpha_2 I + 4\beta_2 \overline{f_z} \overline{f_z}^T \end{bmatrix}. \quad (25)$$

8 LOCALLY INJECTIVE CERTIFICATION

To ensure that the map is locally injective at each iteration, we need to verify that Conditions (4a) and (5) hold (Section 4.1), and backtrack during line search otherwise. As (5) involves infinite number of inequalities: $|f_z(w)| > |\overline{f_z}(w)| \quad \forall w \in \partial\Omega$, we use a simpler sufficient condition based on a finite number of conditions following the approach of [Chen and Weber 2015] which utilizes the fact that the Wirtinger derivatives are Lipschitz continuous. Condition (5) can be enforced on the entire boundary by enforcing it (individually) on many small boundary segments. For each segment

$[v_i, v_{i+1}]$ (see Figure 4 for notations), we compute lower and upper bounds for $|f_z|$ and $|f_{\bar{z}}|$ respectively as follows:

$$|f_z|_{\min} := \frac{|f_z(v_i)| + |f_z(v_{i+1})|}{2} - \frac{L_{f_z} l}{2} \leq \min_{w \in [v_i, v_{i+1}]} |f_z(w)|,$$

$$|f_z|_{\max} := \frac{|f_z(v_i)| + |f_z(v_{i+1})|}{2} + \frac{L_{f_z} l}{2} \geq \max_{w \in [v_i, v_{i+1}]} |f_z(w)|.$$

L_{f_z} and $L_{f_{\bar{z}}}$ are the corresponding Lipschitz constants on the segment, and $l = |v_i - v_{i+1}|$. Condition (5) is then substituted with:

$$|f_z(v_i)| - |f_{\bar{z}}(v_i)| + |f_z(v_{i+1})| - |f_{\bar{z}}(v_{i+1})| \geq (L_{f_z} + L_{f_{\bar{z}}})l, \quad (26)$$

or more concisely:

$$\sigma_2(v_i) + \sigma_2(v_{i+1}) \geq (L_{f_z} + L_{f_{\bar{z}}})l. \quad (27)$$

It is crucial that the sufficient condition above is as tight as possible, as otherwise the Newton step size may become too small and will stop the Newton iterations prematurely. Note that the condition becomes tighter as the Lipschitz constants and/or the length of the segment l decrease. Since denser sampling requires more computations, it is advised to obtain as small as possible Lipschitz constants. In Appendix H, we derive the following Lipschitz constants that can be used on multiply-connected domains. More importantly, our newly-derived constants are significantly smaller than those of BDHM.

$$L_{f_z} = \frac{|f'_z(v_i)| + |f'_z(v_{i+1})|}{2} + \frac{l}{2} \left(\sum_{j=1}^m L_j |s_j - s_{j+1}| + \sum_{j=m+1}^n L_j^h |\varphi_j| \right)$$

$$L_{f_{\bar{z}}} = \frac{|f'_{\bar{z}}(v_i)| + |f'_{\bar{z}}(v_{i+1})|}{2} + \frac{l}{2} \left(\sum_{j=1}^m L_j |t_j - t_{j+1}| + \sum_{j=m+1}^n L_j^h |\psi_j| \right)$$

where $L_j = \frac{1}{2\pi d^2(z_j)}$, $L_j^h = \frac{2}{d^3(\rho_{j-m})}$, $d(z)$ is the distance from a point z to the segment, $s_j = \frac{\varphi_j - \varphi_{j-1}}{z_j - z_{j-1}}$, and $t_j = \frac{\psi_j - \psi_{j-1}}{z_j - z_{j-1}}$. In Figure 7 we compare the Lipschitz constants obtained with our formulas against that of BDHM and show that they are significantly smaller and lead to faster convergence and increased robustness.




Fig. 7. Left: Comparison of the Lipschitz constants of f_z on 529 boundary segments. L_{BDHM} is the approach taken in [Chen and Weber 2015], while L_{ours} is our improved constants. Our constants are smaller on all the segments with an average (number in parentheses above) which is $\times 6.74$ smaller. Right: Comparison of convergence behavior of our solver using BDHM constants (black) vs. ours (red). See how the iterations of the Raptor and Horse deformations stop immaturely when L_{BDHM} is used.

We now turn to Condition (4a) requiring that the boundary integral (or equivalently the number of zeros of f_z) is 0. Theorem

11 in [Chen and Weber 2015] provides a sufficient (but not necessary) condition that implies (4a), albeit their proof assumes that the domain is simply-connected. We show a much stronger result by deriving a new condition which is applicable to multiply-connected domains, and is both necessary and sufficient as long as (27) holds. The following theorem is proved in Appendix C.

THEOREM 8.1. *Let $g(z)$ be an L -Lipschitz continuous holomorphic function in a neighborhood of a simple open curve with length l and two endpoints $v_i, v_{i+1} \in \mathbb{C}$ such that:*

$$|g(v_i)| + |g(v_{i+1})| > Ll, \quad (28)$$

then:

$$\int_{v_i}^{v_{i+1}} \frac{g'(z)}{g(z)} dz = \ln \left| \frac{g(v_{i+1})}{g(v_i)} \right| + i \text{Arg} \frac{g(v_{i+1})}{g(v_i)}, \quad (29)$$

where Arg is the principle branch of the complex argument function.

Theorem 8.1 is applicable to our f_z along any line segment $[v_i, v_{i+1}]$ since our algorithm always verifies that (27) holds (which implies (28)). Consequently we have the following corollary.

COROLLARY 8.2. *Under the assumption that (27) holds, f_z does not vanish inside the multiply-connected polygon P if and only if:*

$$\sum_{j=0}^N \sum_i \text{Arg} \frac{f_z(v_{i+1}^j)}{f_z(v_i^j)} = 0, \quad (30)$$

where the first summation is over the polygonal loops, and the second one is over the segments in each loop.

PROOF. It follows immediately from Cauchy's argument principle and (29) that:

$$2\pi i \mathcal{N} = \oint_{\partial\Omega} \frac{f'_z(w)}{f_z(w)} dw = \sum_{j=0}^N \sum_i \left(\ln \left| \frac{f_z(v_{i+1}^j)}{f_z(v_i^j)} \right| + i \text{Arg} \frac{f_z(v_{i+1}^j)}{f_z(v_i^j)} \right),$$

where \mathcal{N} is the number of zeros of f_z . The $\ln |\cdot|$ terms cancel since $\ln x/y = \ln x - \ln y$ and all the boundary polygons are closed. \square

To conclude, we first verify that (27) holds on all the boundary segments and if so, we simply compute the sum in (30). If it is zero, the map is *guaranteed* to be locally injective everywhere.

9 IMPLEMENTATION

Algorithm 1 provides a pseudocode for our solver. The algorithm can be dramatically accelerated by exploiting parallelism. Most of the steps require only basic dense linear algebra operations and are straightforward to implement. Unlike sparse linear algebra operations, our method greatly benefits when these operations are performed on the GPU. In this section, we point out some of the practical aspects of the implementation. Furthermore, Appendix G contains elaborated description of our GPU implementation. In order to promote reproducibility and to encourage further development of GPU-based solvers, we provide a publicly available reference implementation (<http://github.com/renjie/GILD>).

Algorithm 1 Modified Newton Solver

Input: (ϕ, ψ) ▷ initialization - locally injective map
 $p_i \in \mathcal{P} \subset \Omega, \quad q_i \in \mathbb{C}$ ▷ P2P constraints
Output: (ϕ, ψ) ▷ new locally injective map

```

1: DoPREPROCESSING() ▷ Appendix G
2: loop ▷ Newton iterations
3:    $E_{\text{Def}}^f \leftarrow E^f + \lambda E_{\text{p2p}}^f$  ▷ evaluate energy
4:    $g \leftarrow \nabla E^f + \lambda \nabla E_{\text{p2p}}^f$  ▷ evaluate gradients (G.7),(G.9)
5:    $H \leftarrow \nabla^2 E^f + \lambda \nabla^2 E_{\text{p2p}}^f$  ▷ modified Hessian (G.10)
6:    $(H, g) \leftarrow \text{ELIMINATEVARIABLES}(H, g)$  ▷ Section 5
7:    $Hd = -g, \text{ where } d = \begin{bmatrix} d_\phi \\ d_\psi \end{bmatrix}$  ▷ solve (1) using Cholesky
8:    $(\phi, \psi, t) \leftarrow \text{LINESEARCH}(E_{\text{Def}}^f, d, g, \phi, \psi)$  ▷ backtracking
9:   if  $t \|d\| \leq \epsilon$  then ▷ solver converged
10:    return  $(\phi, \psi)$ 
11:  end if
12: end loop

13: procedure LINESEARCH( $E_{\text{Def}}^f, d, g, \phi, \psi$ ) ▷ Appendix G
14:    $t \leftarrow 1$ 
15:   repeat
16:      $(\bar{t}_z^t, \bar{t}_z^t) \leftarrow \text{COMPUTEWIRTINGER}()$  ▷ Equation (G.11)
17:     if  $E_{\text{Def}}^f(\bar{t}_z^t, \bar{t}_z^t) < E_{\text{Def}}^f + c t (d^T g)$  then ▷  $c \in (0, 1)$ 
18:       if (27) and (30) hold then ▷ locally injective
19:         return  $(\phi + td_\phi, \psi + td_\psi, t)$ 
20:       end if
21:     end if
22:      $t \leftarrow t/2$ 
23:   until MAXNUMSTEPS
24: end procedure

```

9.1 Isometric energy, gradient and Hessian evaluation

While the map and its differentials possess closed-form expressions, we are unable to solve the contour integrals that arise in Equations (12),(13),(14) analytically, and resort to numerical integration. Let $\mathcal{G} = \{w_1, w_2, \dots, w_{|\mathcal{G}|}\} \subset \partial\Omega$ be a set of uniformly distributed samples. We apply the trapezoidal rule to (12) which due to the uniform sampling boils down to a simple sum:

$$E^f \approx \frac{1}{|\mathcal{G}|} \sum_{i=1}^{|\mathcal{G}|} E(w_i). \quad (31)$$

The gradient of E^f is approximated in the same fashion:

$$\nabla E^f \approx \frac{1}{|\mathcal{G}|} \sum_{i=1}^{|\mathcal{G}|} \nabla E(w_i), \quad (32)$$

however, it turns out that for the integral approximation of the Hessian, much coarser approximation can be used with negligible influence on the number of Newton iterations (see Figure 8). Hence, we use an additional set \mathcal{H} of samples which is a subset of \mathcal{G} but




Fig. 8. Hessian approximation. We show the effect of changing the number of Hessian samples $|\mathcal{H}|$ on our algorithm when applied to the depicted deformation of the deer. We observed similar behavior for other models and deformations. With 200 samples or less, we see an increase in the number of iterations needed for convergence (left plot). However, with 400 samples, the convergence rate is identical to that of 50,000 (note the three overlapping curves). On the right plot we see that with 50,000 samples, the algorithm becomes 3 times slower due to the longer time it takes to assemble the Hessian, yet, 400 samples are sufficient, in which case the algorithm converges after 30ms.

contains significantly less samples. The Hessian is approximated:

$$\nabla^2 E^f \approx \frac{1}{|\mathcal{H}|} \sum_{i=1}^{|\mathcal{H}|} \nabla^2 E^+(w_i), \quad (33)$$

where $\nabla^2 E^+(w_i)$ is the *analytic* modification of the pointwise Hessian. We note on a theoretical level, that for any number $|\mathcal{H}|$ (as small as it may be), the Newton step produces a descent direction since the matrix $\nabla^2 E^f$ in (33) is guaranteed to be PSD (strict positive definiteness can be ensured by adding ϵ to the diagonal). In the extreme case that \mathcal{H} is an empty set, the algorithm degrades to gradient descent. We observed that in practice, choosing $|\mathcal{H}| = 0.1 |\mathcal{G}|$ does not degrade the effectiveness of the Newton iterations and provides a good balance between the time it takes to perform the numerical integration and the other steps of the algorithm. We use this setting in all the results. Moreover, since in practice $|\mathcal{H}|$ is much larger than the dimension of our subspace, it is not necessary to add ϵ to the diagonal (Appendix I).

9.2 Line Search

To ensure convergence to a local minimum from an arbitrary starting point we use simple backtracking [Nocedal and Wright 2006, Algorithm 3.1] to find a step size that sufficiently decreases the energy (Algorithm 1, Line 17) and in addition ensures that the new map will be locally injective using the strategy of Section 8 (Algorithm 1, Line 18). This requires evaluations of the energy and the Wirtinger derivatives at different step sizes. Appendix G describes how to compute it efficiently.

10 RESULTS

We use the following default parameters for all our experiments unless stated otherwise. The number of samples is $|\mathcal{G}| = 10^4$, and $|\mathcal{H}| = 0.1 |\mathcal{G}|$. The two energy terms E^f and E_{p2p}^f are balanced using $\lambda = 10^5$. When comparing to mesh-based methods, a mesh with 10,000 vertices is used. User prescribed P2P target positions q_i are visualized as cyan disks. The images of the P2P under the map $f(p_i)$ are visualized as smaller black dots. A black dot which is centered inside a cyan disk indicates P2P satisfaction. Our machine

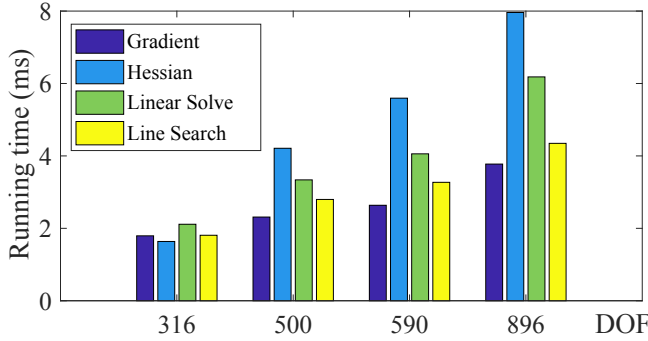


Fig. 9. Algorithm Balance. We show how the runtime of our GPU-based modified-Newton solver is spread across the 4 main stages that compose a single iteration. Each group of bins corresponds to a different model with variable amount of degrees of freedom (variables). For all models we use $|\mathcal{G}| = 10,000$, $|\mathcal{H}| = 1,000$.

runs Windows 10 with Intel Xeon (6 cores) CPU E5-2643 v4 3.40Ghz, 32GB, and has an NVIDIA TitanX graphics card.

In Figure 9 we demonstrate that our solver is well balanced with respect to the four different stages that compose a single Newton iteration: gradient evaluation, Hessian modification-and-assembly, Cholesky linear solve, and line search. Table 2 compares the runtime in milliseconds of a single iteration of various algorithms minimizing the symmetric Dirichlet energy. Among all methods, AQP iterations are the fastest (followed by our GPU iterations), albeit we observe that AQP requires significantly larger number of iterations compared to our GPU/CPU Newton solver. Moreover, the initialization required for AQP (see Section 2) further slows down the runtime (right most column). For comparisons with AQP we run 5 iterations of ARAP local/global [Liu et al. 2008] followed by the projection method of [Kovalsky et al. 2015] as was suggested in [Kovalsky et al. 2016].

Throughout our experiments, we consistently observe that the convergence rate of our algorithm is significantly better than that of the first-order methods (L-BFGS, SLIM, AQP) which we attribute to its second-order nature. Furthermore, the convergence rate is always better even when compared against the Newton-Eigen method which employs a full eigen-factorization [Nocedal and Wright 2006, Equation 3.40] of the Hessian at each iteration, followed by a truncation of negative eigenvalues to $1e-10$. This is demonstrated in Figure 1 (6 vs. 32 iters), Figure 10 (15 vs. 35 iters), Figure 11 (8 vs. 11 iters) suggesting that our Hessian modification is not only faster to compute but also more effective in reducing the energy. Both the eigen-factorization and the solution to the linear system are computed using [cuSOLVER 2017].

The full potential of our deformation algorithm is leveraged when the algorithm runs on graphics hardware, yet our CPU parallel Matlab version (which is approximately 5 times slower on average) is still much faster than the alternatives we compare to. In Figure 10 we show a visual comparison of our algorithm against state-of-the-art methods. For each method we include small images of the maps obtained at intermediate stages of the algorithm to show that for the competing methods, many iterations are needed to get to

Table 2. Comparison of runtime (ms) of a single iteration. SLIM and AQP use a mesh with 10,000 vertices, while the other listed solvers operate within our harmonic subspace with the specified DOF.

	DOF	Ours/GPU	Ours/CPU	Newton Eigen	LBFGS	SLIM	AQP / Initialization
deer	364	8.87	38.8	49.6	54.4	58.3	7.84 1239.7
archery	596	15.31	71.2	95.1	88.3	59.4	9.27 1266.3
giraffe	624	14.05	58.1	86.8	54.9	56.4	8.27 1182.9
rex	392	11.7	61.6	77.6	49.5	57.2	8.52 1141.0
raccoon	320	7.25	34.4	39.9	68.1	53.7	9.01 904.6
raptor	416	10.18	45.2	56.5	82.3	55.1	8.64 883.4

a state where the progress is not visually apparent. Our method is between $\times 25$ to $\times 58$ faster than these competing methods.

We observe that qualitatively, the results obtained with the mesh-based methods tend to concentrate the distortion near the positional constraints. Refining the mesh does not seem to alleviate this issue completely while it increases the computation times even further. We experimentally find that a mesh resolution of 10^4 vertices provides a reasonable balance between quality and speed and use it to evaluate AQP and SLIM. Our algorithm on the other hand is rather indifferent to the mesh resolution and using a mesh with 10^5 vertices, leads to extremely smooth results with no degradation in speed.

Another comparison is given in Figure 11. Our algorithm fully converges in 0.081s (8 iterations) while AQP and SLIM take 2.51s and 9.36s respectively. For the comparisons, we initialize all methods with the identity map which is quite far from the final configuration that is determined by the P2P. In an actual interactive session, the solver utilizes the previous deformation for initialization. Convergence is obtained almost instantly, consequently our algorithm runs in realtime rates.

Figure 7 shows the improvement of our newly-derived Lipschitz constants compared to that of BDHM. These greatly contribute to the success of the algorithm as they ensure that the locally injective validation (line search) has minimal interference, allowing larger Newton steps to be taken.

The appendix contains an additional compilation of high resolution images of deformation results that are obtained with our algorithm. Some of these deformations are intentionally radical and are used to stress test the method and to demonstrate robustness. Our algorithm quickly converges on these extreme deformations and certifies the maps as locally injective. We also include an accompanying video with our paper. Watching it greatly assists in appreciating how fast the algorithm is.

11 SUMMARY AND DISCUSSION

We construct a new finite dimensional subspace of harmonic maps for multiply-connected domains and extend the bounded distortion theorem of [Chen and Weber 2015] to this setting. The harmonic subspace is utilized to devise a novel deformation method with superior performance by casting the locally injective mapping problem as a boundary value problem, minimizing an isometric energy of choice. Within this low dimensional subspace, the induced map and its differentials are C^∞ functions and can be expressed in closed-form,

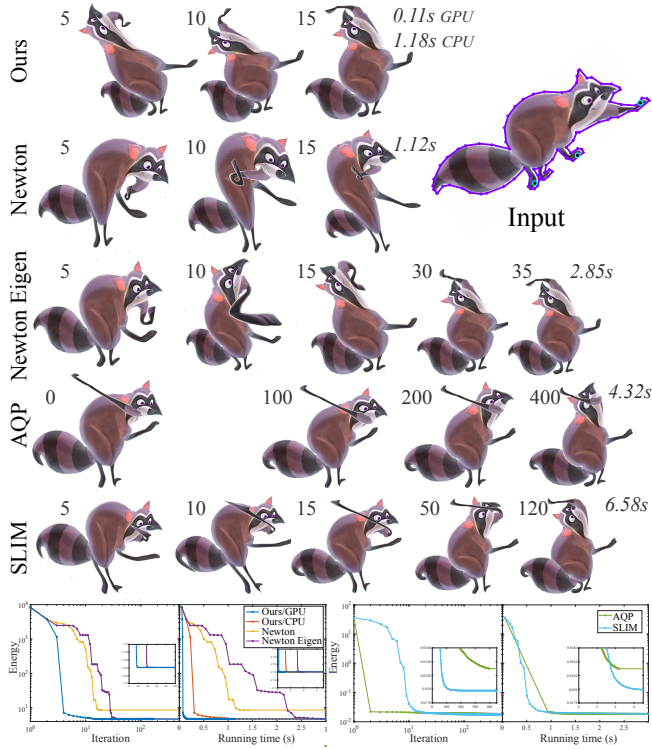


Fig. 10. Visual comparison of the progress of different solvers minimizing $E_{\text{iso}}^f + \lambda E_{\text{p2p}}^f$. The top three rows use the harmonic subspace, while AQP and SLIM are general mesh-based methods. The total number of iterations needed for each method as well as the total runtime is reported on the right of each row. The graphs depict the decrease in energy as a function of iteration and as a function of time, separating harmonic from mesh-based methods as they are scaled differently. The “jump” of 0.9s in AQP’s graph (first iteration) is due to the initialization. Note that due to the large weight λ , the energy at the first few iterations is very large and E_{p2p}^f obscures the behavior of E_{iso}^f . The zoom-in graphs show how the energy decreases more clearly. Our method converges in 15 iterations. Newton’s method halts after 15 iterations due to non-positive definite Hessian. Newton-Eigen performs 35 iterations and is $\times 25$ slower than our GPU solver. AQP and SLIM require 400 and 120 iterations and are $\times 39$ and $\times 58$ slower respectively.

leading to simple, efficient and very accurate evaluation. Moreover, we provide a closed-form recipe for a positive definite modification of the Hessian of these isometric energies that leads to effective-and-efficient Newton iterations. Finally, our algorithm is specifically designed for efficient implementation on modern parallel graphics hardware architectures.

11.1 Limitations and Future Work

The main limitation of our method, which at the same time is also the reason for its superior quality and speed, is its dependence on the reduced dimensional subspace we operate within. Depending on the application, one might desire to produce maps which are not harmonic or not even smooth. Another limitation which is shared by all other methods is the fact that while our method is extremely robust and seems to converge to maps with very low distortion,

global optimality is not guaranteed since the underlying energy is nonconvex. Furthermore, we cannot guarantee that the positional constraints will always be satisfied.

An interesting avenue for future research would be to continue using the harmonic subspace but in a mesh-based setting. While we argue that for the planar deformation setting, the C^∞ basis functions are the most suitable ones, using meshes will allow us to generalize the algorithm to curved manifolds which are the representation of choice for many geometry processing applications, such as surface parameterization, and remeshing. A first step in that direction would be to generalize Theorem 4.2 to smooth curved manifolds, and then to develop a discrete counterpart bounded distortion theorem [Bright et al. 2017].

ACKNOWLEDGMENTS

We thank Edward Chien for illuminating discussions on the integrability of the Wirtinger derivatives on multiply-connected domains and for writing the proof of Theorems 4.1, 4.2, to Markus Steinberger for helping with the GPU implementation, and to Eden Fedida Hefetz for proofreading.

REFERENCES

- L.V. Ahlfors. 1979. *Complex analysis: an introduction to the theory of analytic functions of one complex variable*. McGraw-Hill.
- Noam Aigerman and Yaron Lipman. 2013. Injective and bounded distortion mappings in 3D. *ACM Transactions on Graphics* 32, 4 (2013), 106.
- Noam Aigerman, Roi Poranne, and Yaron Lipman. 2014. Lifted bijections for low distortion surface mappings. *ACM Transactions on Graphics* 33, 4 (2014), 69.
- MOSEK ApS. 2017. *The MOSEK optimization software*. <http://www.mosek.com/>
- Sheldon Axler. 1986. Harmonic functions from a complex analysis viewpoint. *The American mathematical monthly* 93, 4 (1986), 246–258.
- Alon Bright, Edward Chien, and Ofir Weber. 2017. Harmonic Global Parametrization with Rational Holonomy. *ACM Transactions on Graphics* 36, 4, Article 89 (2017).
- Renjie Chen and Ofir Weber. 2015. Bounded distortion harmonic mappings in the plane. *ACM Transactions on Graphics* 34, 4, Article 73 (2015).
- Edward Chien, Renjie Chen, and Ofir Weber. 2016a. Bounded Distortion Harmonic Shape Interpolation. *ACM Transactions on Graphics* 35, 4, Article 105 (2016).
- Edward Chien, Zohar Levi, and Ofir Weber. 2016b. Bounded distortion parametrization in the space of metrics. *ACM Transactions on Graphics* 35, 6, Article 215 (2016).
- cuSOLVER. 2017. CUDA Linear Solvers Library. <http://developer.nvidia.com/cusolver>. (2017). Accessed: 13-May-2017.
- Peter Duren. 2004. *Harmonic mappings in the plane*. Cambridge University Press.
- Xiao-Ming Fu and Yang Liu. 2016. Computing inversion-free mappings by simplex assembly. *ACM Transactions on Graphics* 35, 6 (2016), 216.
- Xiao-Ming Fu, Yang Liu, and Baining Guo. 2015. Computing locally injective mappings by advanced MIPS. *ACM Transactions on Graphics* 34, 4 (2015), 71.
- Eden Fedida Hefetz, Edward Chien, and Ofir Weber. 2017. Fast Planar Harmonic Deformations with Alternating Tangential Projections. *Computer Graphics Forum* 36, 5 (2017), 175–188. Proceedings of Symposium on Geometry Processing 2017.
- Kai Hormann and Michael S. Floater. 2006. Mean value coordinates for arbitrary planar polygons. *ACM Transactions on Graphics* 25, 4 (2006), 1424–1441.
- Takeo Igarashi, Tomer Moscovich, and John F. Hughes. 2005. As-rigid-as-possible shape manipulation. *ACM Transactions on Graphics* 24, 3 (2005), 1134–1141.
- Alec Jacobson, Ilya Baran, Jovan Popović, and Olga Sorkine. 2011. Bounded biharmonic weights for real-time deformation. *ACM Transactions on Graphics* 30, 4, Article 78 (2011).
- Ladislav Kavan, Steven Collins, Jiří Žára, and Carol O’Sullivan. 2007. Skinning with dual quaternions. In *Proceedings of the 2007 symposium on Interactive 3D graphics and games*. ACM, 39–46.
- Shahar Z. Kovalsky, Noam Aigerman, Ronen Basri, and Yaron Lipman. 2015. Large-scale bounded distortion mappings. *ACM Transactions on Graphics* 34, 6, Article 191 (2015).
- Shahar Z. Kovalsky, Meirav Galun, and Yaron Lipman. 2016. Accelerated Quadratic Proxy for Geometric Optimization. *ACM Transactions on Graphics* 35, 4, Article 134 (2016).
- Zohar Levi and Ofir Weber. 2016. On the convexity and feasibility of the bounded distortion harmonic mapping problem. *ACM Transactions on Graphics* 35, 4, Article 106 (2016).




Fig. 11. Raptor. Our solver converges after 8 iterations (0.081s on GPU, 0.31s on CPU) to the depicted map. The Newton-Eigen solver converges to the same result after 11 iterations which took 0.4s due to the longer iteration time. The first-order methods required many more iterations to converge. L-BFGS needs as many as 1,612 iterations to converge to the same result. AQP and SLIM converge to the result depicted on the right after 188 and 170 iterations respectively. Note how the distortion is concentrated around the P2P in this result.

- Yaron Lipman. 2012. Bounded distortion mapping spaces for triangular meshes. *ACM Transactions on Graphics* 31, 4 (2012), 108.
- Yaron Lipman, Johannes Kopf, Daniel Cohen-Or, and David Levin. 2007. GPU-assisted positive mean value coordinates for mesh deformations. In *Symposium on geometry processing*, Vol. 257. 117–123.
- Yaron Lipman, Olga Sorkine, David Levin, and Daniel Cohen-Or. 2005. Linear rotation-invariant coordinates for meshes. *ACM Transactions on Graphics* 24, 3 (2005), 479–487.
- Ligang Liu, Lei Zhang, Yin Xu, Craig Gotsman, and Steven J. Gortler. 2008. A local/global approach to mesh parameterization. *Computer Graphics Forum* 27, 5 (2008), 1495–1504.
- Tiantian Liu, Sofien Bouaziz, and Ladislav Kavan. 2017. Quasi-Newton Methods for Real-Time Simulation of Hyperelastic Materials. *ACM Transactions on Graphics* 36, 3 (2017), 23.
- Tobias Martin, Pushkar Joshi, Miklós Bergou, and Nathan Carr. 2013. Efficient Non-linear Optimization via Multi-scale Gradient Filtering. *Computer Graphics Forum* 32, 6 (2013), 89–100.
- Jorge Nocedal and Stephen Wright. 2006. *Numerical Optimization*. Springer New York.
- Roi Poranne and Yaron Lipman. 2014. Provably good planar mappings. *ACM Transactions on Graphics* 33, 4 (2014), 76.
- Michael Rabinovich, Roi Poranne, Daniele Panozzo, and Olga Sorkine-Hornung. 2017. Scalable Locally Injective Mappings. *ACM Transactions on Graphics* 36, 2, Article 16 (2017).
- John Schreiner, Arul Asirvatham, Emil Praun, and Hugues Hoppe. 2004. Inter-surface mapping. *ACM Transactions on Graphics* 23, 3 (2004), 870–877.
- Christian Schüller, Ladislav Kavan, Daniele Panozzo, and Olga Sorkine-Hornung. 2013. Locally injective mappings. *Computer Graphics Forum* 32, 5 (2013), 125–135.
- Jason Smith and Scott Schaefer. 2015. Bijective Parameterization with Free Boundaries. *ACM Transactions on Graphics* 34, 4, Article 70 (2015).
- Olga Sorkine and Marc Alexa. 2007. As-rigid-as-possible Surface Modeling. In *Proceedings of the Symposium on Geometry Processing*. Eurographics, Switzerland, 109–116.
- Joseph Teran, Eftychios Sifakis, Geoffrey Irving, and Ronald Fedkiw. 2005. Robust quasistatic finite elements and flesh simulation. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer Animation*. ACM, 181–190.
- Amir Vaxman, Christian Müller, and Ofir Weber. 2015. Conformal mesh deformations with Möbius transformations. *ACM Transactions on Graphics* 34, 4, Article 55 (2015).
- Ofir Weber. 2010. *Hybrid Methods for Interactive Shape Manipulation*. Ph.D. Dissertation. Technion - Israel Institute of Technology.
- Ofir Weber. 2017. Planar Shape Deformation. In *Generalized Barycentric Coordinates in Computer Graphics and Computational Mechanics*, Kai Hormann and N. Sukumar (Eds.). CRC Press, Boca Raton, FL, Chapter 7, 28.
- Ofir Weber, Mirela Ben-Chen, and Craig Gotsman. 2009. Complex Barycentric Coordinates with Applications to Planar Shape Deformation. *Computer Graphics Forum* 28, 2 (2009), 587–597.
- Ofir Weber, Mirela Ben-Chen, Craig Gotsman, and Kai Hormann. 2011. A complex view of barycentric mappings. *Computer Graphics Forum* 30, 5 (2011), 1533–1542.
- Ofir Weber and Craig Gotsman. 2010. Controllable conformal maps for shape deformation and interpolation. *ACM Transactions on Graphics* 29, 4, Article 78 (2010).
- Ofir Weber, Olga Sorkine, Yaron Lipman, and Craig Gotsman. 2007. Context-Aware Skeletal Shape Deformation. *Computer Graphics Forum* 26, 3 (2007), 265–274.
- Rhaleb Zayer, Christian Rössl, Zach Karni, and Hans-Peter Seidel. 2005. Harmonic guidance for surface deformation. *Computer Graphics Forum* 24, 3 (2005), 601–609.

APPENDIX[†]

A PROOF OF THEOREM 4.1

We first state the following theorem by Axler [1986]:

THEOREM A.1 (LOGARITHMIC CONJUGATION THEOREM). *Let Ω be a multiply-connected planar domain with N holes K_1, \dots, K_N , and choose N arbitrary points ρ_i in K_i . Then, any real-valued harmonic function $u : \Omega \rightarrow \mathbb{R}$ can be represented as:*

$$u(z) = \operatorname{Re}(g(z)) + \sum_{i=1}^N x_i \ln |z - \rho_i|, \quad (\text{A.1})$$

where g is holomorphic in Ω , and x_1, \dots, x_N are real coefficients.

PROOF. Since $\operatorname{Re}(f)$ and $\operatorname{Im}(f)$ are harmonic in Ω , Theorem A.1 states that there exist holomorphic functions g_1, g_2 and some real constants x_i, y_i such that:

$$\operatorname{Re}(f(z)) = \operatorname{Re}(g_1(z)) + \sum_{j=1}^N x_j \ln |z - \rho_j|,$$

$$\operatorname{Im}(f(z)) = \operatorname{Re}(g_2(z)) + \sum_{j=1}^N y_j \ln |z - \rho_j|.$$

Simple algebraic manipulation gives us:

$$\begin{aligned} f(z) &= \operatorname{Re}(g_1(z)) + i\operatorname{Re}(g_2(z)) + \sum_{j=1}^N (x_j + iy_j) \ln |z - \rho_j| \\ &= \frac{g_1(z) + ig_2(z)}{2} + \frac{\overline{g_1(z) - ig_2(z)}}{2} + \sum_{j=1}^N (x_j + iy_j) \ln |z - \rho_j|, \end{aligned}$$

allowing us to set:

$$\tilde{\Phi} = (g_1 + ig_2)/2, \quad \tilde{\Psi} = (g_1 - ig_2)/2, \quad \omega_i = x_i + iy_i. \quad \square$$

[†]Additional appendices can be found in the supplementary material