Simple Formulas For Quasiconformal Plane Deformations

Y. Lipman
Weizmann Institute
and
V. G. Kim
T. A. Funkhouser
Princeton University

We introduce a simple formula for 4-point planar warping that produces provably good 2D deformations. In contrast to previous work, the new deformations minimizes the *maximum* conformal distortion and spreads the distortion equally across the domain. We derive closed-form formulas for computing the 4-point interpolant and analyze its properties. We further explore applications to 2D shape deformations by building local deformation operators that use Thin-Plate Splines to further deform the 4-point interpolant to satisfy certain boundary conditions. Although this modification no longer has any theoretical guarantees, we demonstrate that, practically, these local operators can be used to create compound deformations with fewer control points and smaller worst-case distortions in comparisons to the state-of-the-art.

Categories and Subject Descriptors:

Additional Key Words and Phrases: Planar warping, Deformation, Conformal, Quasiconformal

1. INTRODUCTION

Planar (2D) warps and deformations are basic operations in image processing with numerous applications, including animation, shape interpolation, registration, media retargeting, image composition and art. Planar deformations are also important in 3D geometric processing for parametrization and intrinsic deformations of surfaces.

Classical methods for planar warping, such as Free-Form Deformations (FFD) [Sederberg and Parry 1986], Thin-Plate Splines (TPS) [Bookstein 1989], and Mean-Value Coordinates (MVC) [Floater 2003] produce warps based on coordinate-wise interpolation and therefore do not have any control over local distortions. Locally they can (and do) introduce arbitrary shears and non-uniform scaling, as shown in Figure 1(b), for example.

Recent 2D warping algorithms have put emphasis on controlling local distortions and thus aim to construct warps that locally preserve angles. Conformal mappings are in this sense optimal as they perfectly preserve angles everywhere. For that reason conformal maps and their approximations have been used extensively for 2D deformations [Igarashi et al. 2005; Lipman et al. 2008; Weber et al. 2009; Weber 2010] and for mesh parameterization [Lévy et al. 2002; Desbrun et al. 2002]. However, conformal mappings have only a small number of degrees of freedom, and cannot, in general, interpolate four or more points and stay injective. For example, Figure 1(d) shows interpolation of four points by Least-Squares Conformal Mapping (LSCM) – note there is a singularity and extreme scaling. For this reason, previous deformation techniques based on conformal mappings had to forsake either interpolation or injectivity: indeed, [Lipman et al. 2008] does not interpolate, the interpolating version of [Weber et al. 2009] is not injective, and [Weber 2010] is locally injective but not interpolatory.

Striving to maintain the local shape preservation of conformal maps while introducing more flexibility, Schaefer et al. [2006] have constructed planar interpolants by locally fitting a similarity using the Moving Least-Squares (MLS) procedure. Still, guarantees or bounds on actually how much conformal distortion is induced are not available. In practice, these MLS maps tend to concentrate the conformal distortion at small areas, often resulting in fold-overs and high conformal distortions, as shown in Figure 1(c).

The goal of this work is to devise *interpolating* 2D warping schemes that have good *conformal distortion* properties while still maintaining properties such as bijectivity and control over scaling. Maps with bounded conformal distortion are called *quasiconformal* [Ahlfors 2006] and recently, researchers have computed such maps for surface registration and parametrization [Zeng et al. 2009; Zeng et al. 2010; Zeng and Gu 2011]. In contrast to previous work, we pose two objectives: 1) we wish to minimize the *maximal* conformal distortion, that is, construct *optimal* quasiconformal maps, and 2) we wish to spread the conformal distortion evenly. As we demonstrate, these objectives lead to deformations that will better preserve local as well-as global properties of shapes.

Although finding optimal quasiconformal map is in general a very hard task, it turns out, surprisingly, that a closed-form solution to this problem can be devised for the particular case of 4 interpolation points, see Figure 1 (a). The solution is given in terms of a very simple formula, defined as composition of two Möbius transformations m_1, m_2 and an affine mapping A:

$$f(z) = m_2 \circ A \circ m_1(z), \tag{1}$$

where $z = x + \mathrm{i}y$ is a complex argument. We will refer to this formula as the 4-Point Interpolant (FPI).

The FPI has several desirable properties: 1) it is defined analytically and easy to apply, 2) it is infinitely smooth and bijection of the plane (possibly with a single point removed), 3) it has *constant* (equally distributed) conformal distortion everywhere (that is the differential of the map has *constant* ratio of maximal to minimal singular values), 4) it minimizes the maximal conformal distortion over all possible mappings of a certain class, 5) it has an analytic inverse with the *same* conformal distortion as the forward mapping everywhere, and 6) it has closed-form formulas for computing m_1, m_2, A for any given two sets of four points.

Finding the optimal quasiconformal map for more than 4 interpolation points is, unfortunately, much harder problem and we do not provide a solution to that problem in this paper. Nevertheless, in this paper we demonstrate how the 4-point formula (FPI) can be practically used as an approximate solution to a more general class of deformation operators that satisfy some extra boundary conditions. In particular, we use the FPI scheme repeatedly as a basic building block for constructing simple and effective deformation operators that are comparable to state of the art deformation algorithms in terms of deformation quality, simplicity of the algorithm, and the amount of input required from the user to guide the deformation.

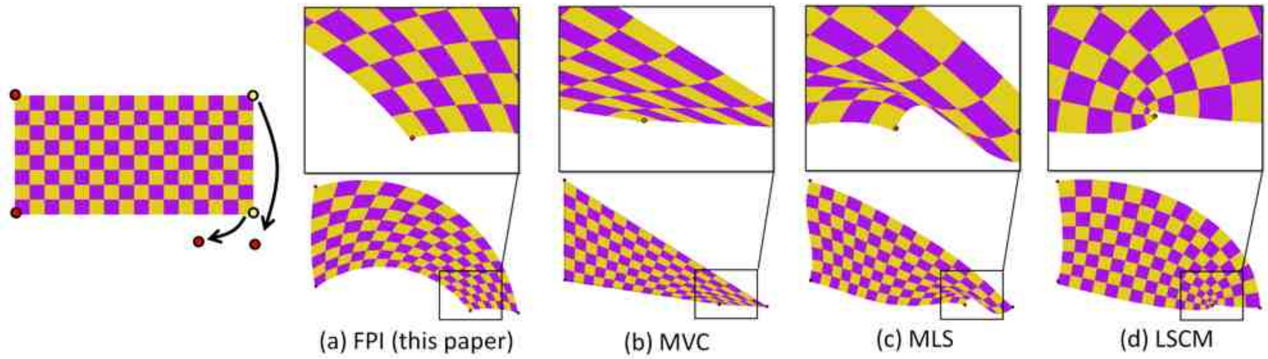

Fig. 1: *Deformation of a rectangle domain based on four interpolation points placed at the corners (left). The results of four methods are shown (left to right): FPI (this paper), MVC [Floater 2003], MLS [Schaefer et al. 2006], LSCM [Lévy et al. 2002] and [Igarashi et al. 2005].*

2. 4-POINT WARPING

In this section, we present the key ingredient of this paper: the 4-point interpolant (FPI) formula. Our goal is to answer the following question: given an ordered set of four source points $Z = \{z_1, z_2, z_3, z_4\} \subset \mathbb{C}$, where $\mathbb{C} = \{x + iy \mid x, y \in \mathbb{R}\}$ denotes the complex plane, and four target points $W = \{w_1, w_2, w_3, w_4\} \subset \mathbb{C}$, what is the "most conformal" way to interpolate these points with a bijective map of the plane? We will show that under certain assumptions the FPI minimizes the maximal conformal distortion and therefore is optimal in the L^∞ sense. We will construct formulas to find m_1, m_2, and A for arbitrary quadruplets Z, W. In the next section we will prove, among other properties, that the FPI also spreads the conformal distortion equally everywhere. We will assume, without limiting our discussion, that the quadruplets Z and W are bounding a four sided polygon and that they are ordered in counter-clockwise fashion (different order will lead to a different map).

2.1 A simple case: parallelograms

In this subsection we present a solution to the "most conformal" mapping problem in the restrictive case that both point sets, Z, W, consist of corners of two parallelogram, $P(\nu, \xi)$ and $P(\widehat{\nu}, \widehat{\xi})$ (resp.), where by $P(\nu, \xi)$ we denote the interior of a parallelogram with corners $\{0, \nu, \nu + \xi, \xi\}$ (we can always translate one corner to the origin).

When looking for an optimal map, one should define a collection of maps to search in; we want to consider a family of mappings $\mathcal{F} = \{f\}$, from which we search for the optimal $f^* \in \mathcal{F}$, that is a map that minimizes the maximal conformal distortion, where conformal distortion is defined at every point as the ratio of the maximal to minimal singular values of the differential of the map (aspect ratio of the ellipse). Since we want our map to be defined on the entire plane, it is natural to think of "tileable" or "periodic" maps. Given two parallelograms $P(\nu, \xi)$ and $P(\widetilde{\nu}, \widetilde{\xi})$, a periodic map is defined by the rule

$$f(z + m\nu + n\xi) = f(z) + m\widehat{\nu} + n\widehat{\xi},$$

where $m, n \in \mathbb{Z}$ (integers). Intuitively, we simply require that the map f is tileable over the lattice defined by the parallelograms, see Figure 2 (a). Furthermore, we will require that f is differentiable across the boundaries of the parallelogram. Another way to think about these periodic maps is by "stitching" the two opposite sides of the parallelograms and considering differentiable maps between the two resulting tori.

In this huge collection of periodic maps, there is one special map that minimizes the maximal conformal distortion. Interestingly, it is a very simple map: the affine map that takes $P(\nu, \xi)$ to $P(\widetilde{\nu}, \widetilde{\xi})$. In the appendix, based on arguments due to Ahlfors [Ahlfors 2006], we prove that every differentiable periodic map $f : P(\nu, \xi) \to P(\widetilde{\nu}, \widetilde{\xi})$ must have the following lower bound on the maximal conformal distortion, denoted here by K_f:

$$K_f \geq e^{d_H\left(\mathrm{Im}(\nu/\xi), \mathrm{Im}\left(\widetilde{\nu}/\widetilde{\xi}\right)\right)}, \tag{2}$$

where $d_H(z, w) = \log\left[\frac{|w-\bar{z}|+|w-z|}{|w-\bar{z}|-|w-z|}\right]$ is the hyperbolic distance in the upper half-plane. It is not hard to check (and is also shown in the Appendix) that the affine map taking $P(\nu, \xi)$ to $P(\widetilde{\nu}, \widetilde{\xi})$ achieves this bound and is therefore optimal.

This observation provides a direct way to produce an interpolatory and bijective map minimizing the maximal conformal error for four control points – simply use the affine map defined as:

$$A(z) = w_1 + \ell_1(z - z_1) + \ell_2\overline{(z - z_1)}, \tag{3}$$

where ℓ_1, ℓ_2 specify the linear transformation $L(z) = \ell_1 z + \ell_2 \bar{z}$ on a complex point (z) determined by solving the following 2×2 linear system:

$$\begin{pmatrix} (z_2 - z_1) & (\bar{z}_2 - \bar{z}_1) \\ (z_3 - z_1) & (\bar{z}_3 - \bar{z}_1) \end{pmatrix} \begin{pmatrix} \ell_1 \\ \ell_2 \end{pmatrix} = \begin{pmatrix} w_2 - w_1 \\ w_3 - w_1 \end{pmatrix}. \tag{4}$$

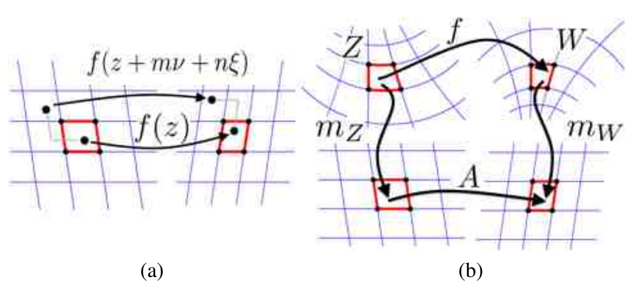

Fig. 2: *Construction of the FPI, see the text for details.*

2.2 The general case: quadrilateral

In this subsection, we will present a general solution to the "most conformal" mapping problem for 4 point interpolants, that is, we consider the case where Z, W are two general planar quadruplets (counter-clockwise ordered) and ask how to interpolate this data while being as conformal as possible in the maximum norm sense.

The key insight that allows us to use the simple solution for parallelograms presented in the previous subsection for general two quadruplets Z, W, is the observation that, from the *conformal* point of view, any quadruplet of points can be seen as corners of some *circular* parallelogram.

To understand this statement and how we use it to solve the problem stated above, let us first define, for any ordered quadruplet $Z = \{z_1, z_2, z_3, z_4\}$ (we will do similarly for W), circular edges that will turn Z into a "parallelogram".

PROPOSITION 2.1. *Given a quadruplet* $Z = \{z_1, z_2, z_3, z_4\} \subset \mathbb{C}$ *(prescribed in counter-clockwise order) there exists a unique fifth point* z_∞ *such that the following conditions hold:*

(1) *There are four circles (where straight lines are considered as circles with infinite radii) defined by this fifth point and every consecutive pair of points* $z_i\, z_{i+1}$. *These four circles define four circular edges that form the circular parallelogram.*

(2) *Each pair of opposite arcs define two circles that meet only at this fifth point (osculant circles).*

(3) *This extra fifth point is in the exterior of the circular parallelogram ("outside" is defined using the order of the points),*

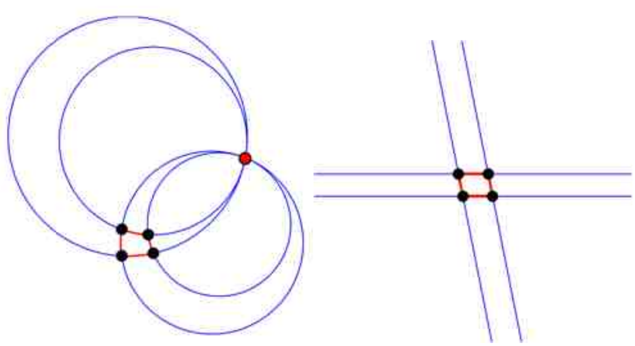

(a) (b)

Fig. 3: *The circular parallelogram shown in red (a), and its Euclidean counterpart (b).*

Figure 3(a) shows an example where the points Z are shown as black disks and the unique fifth point is shown in red.

Before we explain how to use the observation or prove it, let us first explain why we call such circular-edged quadrilateral a circular parallelogram: using a conformal bijective map of the extended plane (the complex plane added with infinity as a legitimate point), one can map this circular edged quadrilateral to a standard Euclidean parallelogram. Indeed, taking the fifth point to ∞ via a Möbius transformation (to be defined) will leave two pairs of parallel lines (that only meet at infinity) forming the four corners of the Euclidean parallelogram (see Figure 3(b) where we did exactly that for the example in (a)). In particular, opposite angles in the circular parallelogram are equal, a characterizing property for Euclidean parallelograms. To prove Proposition 2.1 for every ordered quadruplet, we will prove an equivalent statement:

PROPOSITION 2.2. *Given a quadruplet* $Z = \{z_1, z_2, z_3, z_4\} \subset \mathbb{C}$ *(prescribed in counter-clockwise order) there exists a unique (up to a similarity transformation) Möbius transformation* m_Z *that takes* Z *to corners of a parallelogram* P_Z*, while preserving the orientation of the boundary points.*

Where Möbius transformations are defined by the formula

$$m(z) = \frac{az+b}{cz+d}, \quad ad - bc \neq 0, \quad a, b, c, d \in \mathbb{C}, \tag{5}$$

and constitute the group of conformal maps bijectively mapping the extended complex plane onto itself.

Let us explain why Propositions 2.1 and 2.2 are equivalent.

LEMMA 2.3. *Proposition 2.1 and Proposition 2.2 are equivalent.*

PROOF. First, assuming Proposition 2.2 is true, we can define $z_\infty = m_Z^{-1}(\infty)$ and the respective inverse image of the two pairs of straight lines forming the parallelogram will provide the desired circular parallelogram. In the other direction, assuming Proposition 2.1, we can define m_Z to be any Möbius transformation such that $m_Z(z_\infty) = \infty$. The conditions on the four circles forming the circular parallelogram will assure that their image under m_Z consists of two pairs of parallel lines. The uniqueness in both cases is clear. □

Next, we explain how the above observations are useful to solve the problem stated above. Since a Möbius transformation is a bijective conformal map, using it to map a quadruplet to parallelogram's corners does not introduce any conformal distortion and reduces the general problem back to parallelograms, as follows. Consider two general quadruplets, Z, W, and denote by m_Z the Möbius transformation taking Z to a parallelogram's corners P_Z, and m_W the Möbius transformation taking W to parallelogram's corners P_W. Furthermore, let A be the affine map taking the corners of one parallelogram P_Z to corners of another parallelogram P_W, as defined in eq. (3). Then our final 4-point interpolant is defined as composition of m_Z, A, and m_W^{-1} (see Figure 2 (b)), that is,

$$f(z) = m_W^{-1} \circ A \circ m_Z(z), \tag{6}$$

where the inverse of a Möbius transformation is also a Möbius transformation and is calculated by simply inverting the 2×2 coefficient matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$.

The idea is that since Möbius transformations are conformal, they do not introduce any conformal distortion, and therefore the 4-point interpolant $f(z)$, interpolating the two quadruplets Z, W, have the same conformal distortion as the affine map A, which is known to be optimal.

The formulas for finding a Möbius transformation m_Z mapping a quadruplet to a parallelogram P_Z are summarized in Algorithm 1 below. The pseudocode for calculating the FPI's different components (m_Z, m_W, A) is provided in Algorithm 2.

2.3 Derivations of formulas for Möbius mapping to a parallelogram (Proof of Proposition 2.2)

Let us now derive closed-form formulas for finding m_Z (m_W will be computed similarly), in doing so we will prove Proposition 2.2: the proof will outline explicit formulas for finding m_Z given Z.

Algorithm 1: quadruplet_to_parallelogram (Z)

Input: Source points $Z = \{z_1, z_2, z_3, z_4\}$
Output: Möbius transformation $m = \frac{az+b}{cz+d}$
and a linear map $L(z) = z + \ell\overline{z}$

$G = \left\{g_j = \exp\left(\mathrm{i}\frac{2\pi j}{n}\right)\right\}_{j=0,..,3}$
$M = \left[Z \mid \mathbf{1} \mid -ZG \mid -G \mid -Z\overline{G} \mid -\overline{G}\right]$
$USV^* = SVD(DM)$
$u = V(:,3)\ ,\ v = V(:,4)$
`/* Solve the quadratic equation in` $t \in \mathbb{C}$ `*/`
$t^2(v_6v_3 - v_4v_5) + t(v_6u_3 + u_6v_3 - u_5v_4 - v_5u_4) +$
$(u_6u_3 - u_5u_4) = 0$
$x = u + t_1 v\ ;\ y = u + t_2 v$
`/* Two candidate solutions */`
$a = x_1\ ;\ b = x_2\ ;\ c = x_3\ ;\ d = x_4\ ;\ \ell = x_6/x_4$
$a = y_1\ ;\ b = y_2\ ;\ c = y_3\ ;\ d = y_4\ ;\ \ell = y_6/y_4$
Return the solution with the smaller $|\ell|$.

Algorithm 2: FPI(Z,W)

Input: Source points $Z = \{z_1, z_2, z_3, z_4\}$,
Target points $W = \{w_1, w_2, w_3, w_4\}$
Output: FPI transformation $f(z) = m_W^{-1} \circ A \circ m_Z$

m_Z = quadruplet_to_parallelogram(Z)
m_W = quadruplet_to_parallelogram(W)
A = calculate_affine_map($m_Z(Z)$,$m_W(W)$)
Return $f = m_W^{-1} \circ A \circ m_Z$

First, we note that the problem formulated in the proposition can be described as follows: given a quadruplet Z (ordered in counter clockwise fashion) we look for a Möbius transformation m_Z and an invertible, orientation preserving, linear mapping L, such that

$$m_Z(z_j) = L(g_j),\ \ j = 1..4, \tag{7}$$

where $G = \left\{g_j = \exp\left(\frac{\mathrm{i}2\pi(j-1)}{4}\right), j = 1..4\right\}$, are four corners of a square. To solve this equation we plug the general expression for a Möbius transformation (5), and a linear map $L(z) = \ell_1 z + \ell_2\overline{z}$. However, since we can assume $\ell_1 \neq 0$ (since otherwise we get an orientation-reversing linear map), we can scale both sides of eq. (7) by $1/\ell_1$. So it is enough to consider $L(z) = z + \ell\overline{z}$ for the linear part:

$$\frac{az_j + b}{cz_j + d} = g_j + \ell\overline{g}_j,\ \ j = 1..4. \tag{8}$$

Multiplying both sides by $cz_j + d$ and rearranging we get the following system of 4 nonlinear equations in 5 unknowns (written in matrix form):

$$\left[Z \mid \mathbf{1} \mid -ZG \mid -G \mid -Z\overline{G} \mid -\overline{G}\right](a, b, c, d, c\ell, d\ell)^T = \mathbf{0}, \tag{9}$$

where we denote (with a slight abuse of previous notation) $Z, G \in \mathbb{C}^{4\times 1}$ to be column vectors of 4-complex points $(z_1, .., z_4)^T, (g_1, .., g_4)^T$ (respectively), $ZG \in \mathbb{C}^{4\times 1}$ denotes their coordinate-wise multiplication, and $\mathbf{1}, \mathbf{0} \in \mathbb{C}^{4\times 1}$ the column vector of ones and zeros (respectively). Denote the matrix in eq. (9) by $M \in \mathbb{C}^{4\times 6}$. In the generic case the rank of M is exactly 4 (since the columns are samples of linearly independent polynomials). Next, perform the Singular Value Decomposition

$$M = USV^*,$$

where $U \in \mathbb{C}^{4\times 4}, V \in \mathbb{C}^{6\times 6}$ unitary matrices, superscript $*$ represents the conjugate transpose, and $S \in \mathbb{C}^{4\times 6}$ diagonal matrix with the singular values $\sigma_1 \geq \sigma_2 \geq ... \geq \sigma_4 > 0$ along its diagonal. For a solution of the form $x = (a, b, c, d, c\ell, d\ell)^T$ to exist, x should satisfy the relation

$$x(6)/x(4) = x(5)/x(3). \tag{10}$$

The two least-significant (i.e., corresponding to smallest singular values) right singular vectors u, v (columns of V) have zero singular values. Since x and λx (λ is any complex number) result in the same solution (as a Möbius transformation is set up-to a multiplicative constant) we can search a solution of the form $x = u + tv$. Enforcing relation (10) on x we get a quadratic equation in t (over $\mathbb{C}$) with two roots $t_1, t_2 \in \mathbb{C}$. Both $x_1 = u + t_1 v$, $x_2 = u + t_2 v$ satisfy system (9) and eq. (10), and therefore solve the problem.

Let us show that the two solutions, x_1, x_2, correspond to two distinct Möbius transformations m_+, m_-, and furthermore that only one of them, which we will denote w.l.o.g by m_+, does not flip inside-out the interior of the polygon Z. First, let us show that the two solutions are distinct and characterize how they relate to one another. Take x_1 and set a Möbius transformation m based on its first four coordinates. Then $m(Z)$ is a parallelogram P_Z with its center (intersection of diagonals) placed at the origin. Now, let us apply the Möbius transformation $\widetilde{m}(z) = 1/z$ to the parallelogram P_Z. From the symmetry of the parallelogram w.r.t the origin, we see that $\widetilde{m}(P_Z)$ are also corners of some parallelogram that is centered at the origin. Since Möbius transformations form a group, composing m with $\widetilde{m}$ results in a second solution m^*. Note that the order of the boundary points of $\widetilde{m}(P_Z)$ is now flipped. Therefore, only one of the Möbius transformations preserves the orientation of the boundary and that is the desired Möbius transformation. Since the Jacobian of the linear map L can be written in complex notation as $J_L = 1 - |\ell|^2$, we can find the good solution by taking the solution x_1 or x_2 that results in $|\ell| < 1$ (the smaller among the two solutions). In non-generic situations $x = v$ could be a solution to the system ($t = \infty$), in that case we get a linear equation in t and we still end up with exactly two solutions to (8) where only one of them is the correct solution. This constructive proof suggests Algorithm 1 that is very simple and requires only one matrix singular value decomposition.

3. THE PROPERTIES OF THE FPI.

In this section we describe the main properties of the FPI. Since the FPI has a very simple analytic formula, it has properties that are easy to prove.

3.1 Smooth bijection of the punctured plane

PROPERTY 1. *The FPI $f = m_W^{-1} \circ A \circ m_Z$ is a C^∞ bijective map $f : \mathbb{C} \setminus z_\infty \to \mathbb{C} \setminus w_\infty$ (punctured planes), where z_∞ is defined by $z_\infty = m_Z^{-1} \circ A^{-1} \circ m_W(\infty)$, and w_∞ is defined by $w_\infty = m_W^{-1} \circ A \circ m_Z(\infty)$.*

PROOF. f is a composition of bijective C^∞ maps from the extended complex plane to itself, therefore, it is bijective C^∞ from the standard complex plane, possibly with one point removed (the one that is mapped to ∞), to the complex plane, again with possibly one point removed (the image of ∞). Figuring out the image and inverse image of ∞ leads to the above equations specified for z_∞ and w_∞. □

3.2 Constant conformal distortion

PROPERTY 2. *The FPI f has constant conformal distortion everywhere.*

We will use standard complex-theory notations (see e.g., [Ahlfors 2006] page 3). Briefly, $z = x + \mathrm{i}y \in \mathbb{C}$ will denote the complex argument and the complex differentials and derivatives are defined by $dz = dx + \mathrm{i}dy$, $d\overline{z} = dx - \mathrm{i}dy$, and $\partial_z = \partial_x - \mathrm{i}\partial_y$, $\partial_{\overline{z}} = \partial_x + \mathrm{i}\partial_y$, respectively. The differential of a complex valued function $f : \mathbb{C} \to \mathbb{C}$ using this notation is $df = f_z dz + f_{\overline{z}} d\overline{z}$. The benefit in this representation in our context is that the Cauchy-Riemmann equations are simply $f_{\overline{z}} = 0$. A common measure of conformal distortion is then

$$D_f = \frac{|f_z| + |f_{\overline{z}}|}{|f_z| - |f_{\overline{z}}|} \geq 1, \tag{11}$$

and D_f equals one if and only if f is conformal. For orientation preserving maps D_f can be shown to be the ratio between the maximal and minimal singular values of df (for orientation reversing one gets the negative ratio).

PROOF. Calculating the conformal distortion of (6) using eq.(11) and the standard product rule for complex derivatives (see e.g., [Ahlfors 2006]) leads to

$$D_{m_W^{-1} \circ A \circ m_Z}(z) = \frac{|\ell_1| + |\ell_2|}{|\ell_1| - |\ell_2|}.$$

This shows that the conformal distortion of the FPI is constant everywhere and equals the conformal distortion of the affine map between the corresponding parallelograms. □

3.3 Minimal maximum conformal distortion

PROPERTY 3. *The FPI minimizes the maximal conformal distortion*

$$f = m_W^{-1} \circ A \circ m_Z = \underset{\widetilde{f} \in \widetilde{\mathcal{F}}}{\operatorname{argmin}} \max_z D_{\widetilde{f}}(z),$$

among a family $\widetilde{\mathcal{F}} = \left\{ \widetilde{f} \right\}$ of periodic mappings that take one quadruplet Z to the other W.

The mapping collection $\widetilde{\mathcal{F}}$ that we are considering consists of the entire collection of differentiable bijective mappings $\widetilde{f}$ that map the (unique) torus defined by Z to the (unique) torus defined by W, while interpolating the corners $\widetilde{f}(z_i) = w_i$, $i = 1..4$; where the torus defined by Z (similarly for W) is the unique circular-edged quadrilateral, the existence of which is guaranteed by Proposition 2.1, where each pair of opposite edges are identified with a Möbius transformation rather than just a translation like the Euclidean case. In that sense every quadruplet can be seen as a torus, and a differentiable periodic mapping is a map that is well-defined on this torus (behave consistently across circular boundaries).

Among all differentiable mappings that satisfy these periodic boundary conditions the FPI minimizes the maximal conformal distortion. For example, in Figure 4 we show a comparison of the FPI to another map $\widetilde{f}_{mvc} \in \mathcal{F}$ that is achieved by using Mean Value Coordinates to interpolate the prescribed boundary conditions. Note that, as expected, the $\widetilde{f}_{mvc}$ has higher maximal conformal distortion.

PROOF. Let us denote by m_Z (m_W) the Möbius transformation taking Z (W) to corners of some standard Euclidean parallelogram

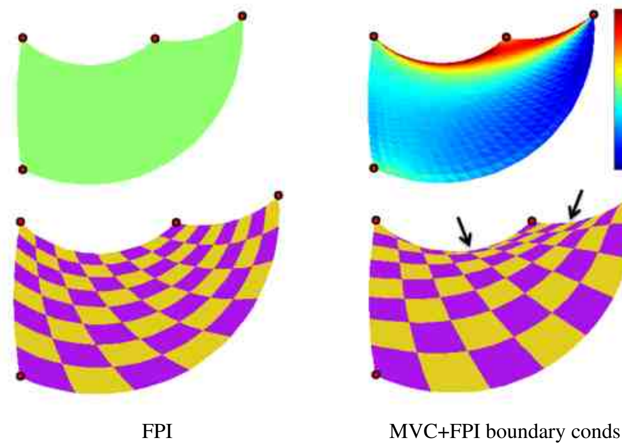

Fig. 4: *We compare the FPI and MVC where we set the boundary behavior to match the FPI boundary behavior. As our theoretical analysis shows indeed the FPI achieves smaller maximal conformal distortion (conformal distortion is depicted in top row where dark blue is zero distortion and dark red is high distortion). The MVC tends to distribute the conformal distortion unevenly, and in this case even cause fold-overs (see marked area).*

P_Z (P_W). Every periodic map between the two parallelograms $f \in \mathcal{F} : P_Z \mapsto P_W$ (defined in Section 2.1) can be converted to a map $\widetilde{f} \in \widetilde{\mathcal{F}}$ by the simple rule $\widetilde{f} = m_W^{-1} \circ f \circ m_Z$. Note that $\widetilde{f}$ and f have the same conformal distortion (since m_Z, m_W are conformal) and that this procedure provides a bijection between $\mathcal{F}$ and $\widetilde{\mathcal{F}}$. Therefore,

$$\begin{aligned}
\widetilde{f} &= \underset{g \in \widetilde{\mathcal{F}}}{\operatorname{argmin}} \max_z D_g(z) \\
&= m_W^{-1} \circ \left[\underset{g \in \mathcal{F}}{\operatorname{argmin}} \max_z D_{m_W^{-1} \circ g \circ m_Z}(z) \right] \circ m_Z \\
&= m_W^{-1} \circ \left[\underset{g \in \mathcal{F}}{\operatorname{argmin}} \max_z D_g(z) \right] \circ m_Z \\
&= m_W^{-1} \circ A \circ m_Z,
\end{aligned}$$

where the last equality is due to the optimality of the affine map between parallelograms, as explained in Section 2.1. □

3.4 Inverse map

PROPERTY 4. *The inverse of the FPI is simply $f^{-1}(w) = m_Z^{-1} \circ A^{-1} \circ m_W(w)$, and therefore is also an FPI.*

The proof is obvious. Note that the inverse FPI f^{-1} is precisely the FPI that we would get if we were to solve the reverse problem $W \to Z$. Even more interesting is the fact that the conformal distortion of the inverse map equals the conformal distortion of the forward map, $D_f = D_{f^{-1}}$ (verified with a direct computation). Note that this "symmetric" property is a unique outcome of the FPI construction and does not exist, as far as we are aware, in other methods.

3.5 Alternative solution

Let us conclude this section by reviewing an alternative solution for the 4-point mapping problem by considering a *different family* of mappings $\widetilde{\mathcal{F}}$, namely the collection of bijective and differentiable maps mapping the interior of the quadrilateral defined by Z (with straight edges) to the interior of the quadrilateral defined by W. In this case the optimal solution that minimizes the maximal conformal distortion can be constructed as follows: first map each quadrilateral to a rectangle conformally, and then stretch one rectangle onto the other. It is possible to prove the optimality of this solution w.r.t the space $\widetilde{\mathcal{F}}$ described above (e.g., see [Ahlfors 2006] page 6). However, this solution has several drawbacks: first, the map cannot (generally) be extended outside the parallelograms.

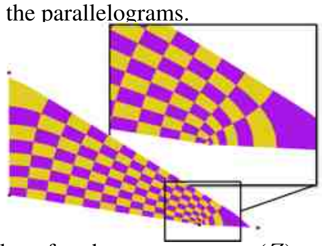

Second, the family of mappings considered for this solution satisfies stricter boundary conditions; the mappings preserve the straight boundary edges of the source and target quadrilateral. The inset shows the result of the above procedure for the same source (Z) and target (W) points as Figure 1. Note however, that the conformal distortion is higher than the FPI result $(1.86 > 1.79)$ and the maximal area distortion is considerably higher $(12.57 > 2.39)$. Lastly, the conformal mapping of a quadrilateral to a rectangle needs to be numerically approximated and will render the solution slower to compute.

4. LOCAL FPI FOR CONSTRAINED DEFORMATIONS

In this section we investigate how the FPI scheme can be used to create elaborate deformations of 2D domains.

The key idea is to create deformation operators with local support that are as similar as possible to the FPI. Although for this case we do not have any theoretical guarantee that our solution minimizes the maximal conformal distortion nor that it approximates such an optimal solution, we show that, practically, the FPI provides a good basis/approximation for such deformations.

Furthermore, we believe that four points are the intuitive number of control points for a human to manipulate simultaneously - e.g., for a deformation application on a touch screen.

4.1 User interface

We will use the FPI *locally*, such that the deformation is performed "inside" a user's defined Region Of Interest (ROI), while connecting smoothly to the "outside" part where we perform a constant similarity transformation (e.g., the identity).

As an example, in this section, we will construct two deformation operators in this spirit. Later, in Section 5, we demonstrate that, together, these operators can create a wide range of deformations competitive with previous work in terms of quality of the deformations, simplicity of the algorithm, and in the number of user handles used to guide the deformation.

For the rest of this section we denote by $\Omega \subset \mathbb{C}$ the domain we wish to deform, and we assume that Ω is simply connected, where simply connected means that every loop can be continuously contracted to a point without leaving Ω.

We will construct two types of deformation operators: as shown in Figure 5, the user clicks on four points $Z = \{z_1, z_2, z_3, z_4\}$ (green disks) and chooses two edges $e^Z_\alpha = z_\alpha\, z_{\alpha+1}, e^Z_\beta = z_\beta\, z_{\beta+1}$ (colored in blue) defining the ROI (region 1). The user then moves the "free" vertices (each marked with four arrows), prescribing new locations W of the initial four points Z. Let us further denote the deformed edges by $e^W_\alpha = w_\alpha\, w_{\alpha+1}$, and $e^W_\beta = w_\beta\, w_{\beta+1}$.

The deformation of the ROI (region 1) is done using FPI, while the deformation of the "outer" regions (regions 2,3) is defined as the unique constant similarity transformation defined by the transformation of the edges e^Z_α, e^Z_β. In case the two chosen edges are adjacent (Figure 5(b)) they assumed to undergo the same similarity (e.g., the identity mapping).

At this point the deformations of the outer region(s) is set by the edges e^Z_α, e^Z_β and has zero conformal distortion, and in case the length of the edges is preserved, by a perfect isometry (rigid motion).

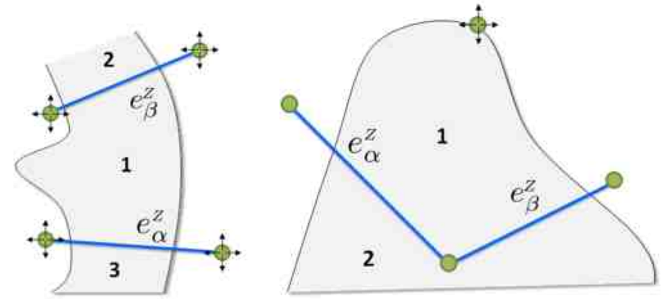

(a) (b)

Fig. 5: *Two types of operators implemented.*

4.2 Constrained FPI

We are left with the main part of deforming the ROI with as-low-as-possible distortion while smoothly connecting to the similarities at the edges e^Z_α, e^Z_β. This means that, in the spirit of previous sections, we are facing the following problem: we are given two quadruplets, Z and W, and we wish to find a map f that minimizes the maximal conformal distortion among all the maps that interpolate these quadruplets of points $f(z_j) = w_j$, $j = 1..4$, and furthermore, interpolate the values and derivatives of the similarities along the two prescribed edges e^Z_α, e^Z_β.

This problem is slightly different from the problem solved by the original FPI introduced in previous sections and it is unlikely that a closed form solution to this problem can be found. As a matter of fact, even trying to numerically approximate this map seems challenging (mainly because of the min-max-norm formulation and the huge size of the space of possible maps).

Nevertheless, as we demonstrate next, the FPI can be used to devise an approximate solution. Using notations from Section 2, we can think of the Möbius transformations m_Z, m_W that takes Z, W to parallelograms P_Z, P_W (resp.) as change of coordinates. In these new coordinates the FPI is a simple affine map. In the current case, after the change of coordinates, we have extra constraints along two edges (now transformed by m_Z, m_W). Hence, instead of simple affine map (which is optimal), we will look for a map φ that is *closest* to affine and satisfy these extra constraints.

Measuring a "distance" between a C^2 map φ to an affine map (denote by Aff the planar affine group) can be done using the well known second-order Sobolev semi-norm:

$$dist(\varphi, Aff) := \|\varphi\|_{W^{2,p}} = \left(\|\varphi_{xx}\|^p_{L_p} + 2\, \|\varphi_{xy}\|^p_{L_p} + \|\varphi_{yy}\|^p_{L_p} \right)^{1/p},$$

where $\|\cdot\|_{L_p}$ denotes the $L_p = L_p(\Omega)$ norm in Ω where $1 \leq p \leq \infty$. To achieve closed form solution in this case we will pick $p = 2$, and $\Omega = \mathbb{C}$ to get the well-known Thin-Plate Spline (TPS) energy [Wendland 2005]:

$$\varphi = \underset{\widetilde{\varphi}}{\operatorname{argmin}} \int_{\mathbb{C}} \left[\left| \frac{\partial^2 \widetilde{\varphi}}{\partial x^2} \right|^2 + 2 \left| \frac{\partial^2 \widetilde{\varphi}}{\partial x \, \partial y} \right|^2 + \left| \frac{\partial^2 \widetilde{\varphi}}{\partial y^2} \right|^2 \right] dx \, dy \,,$$

the minimizers of which are the Thin-Plate Splines.

Motivating by this observation, our plan is to define the mapping of the ROI via the map

$$f = m_W^{-1} \circ \varphi \circ m_Z, \tag{12}$$

where φ is a TPS function

$$\varphi(z) = \sum_{j=1}^{J} b_j \phi(|z - c_j|) + A(z), \tag{13}$$

where $\phi(r) = r^2 \log(r)$, $\{c_j\} \subset \mathbb{C}$ are the interpolation centers and $\{b_j\} \subset \mathbb{C}$, $A(z) = \ell_1 \, z + \ell_2 \, \overline{z} + \ell_3$, $\{\ell_k\} \subset \mathbb{C}$ are coefficients and an affine map (resp.) to be set for satisfying a set of interpolation constraints

$$\varphi(c_j) = d_j \,, j = 1, .., J \,, \tag{14}$$

where $\{d_j\} \subset \mathbb{C}$ are positional constraints.

Intuitively, φ can be seen as the most affine map (in the sense of minimizing its second derivatives' L_2 norm) that satisfy the constraints (14); in case we do not pose any edge constraints, φ would be an affine map and therefore reproduce the FPI.

Calculating $\{b_j\}$, $\{\ell_k\}$ given the interpolation centers $\{c_j\}$ and constraints (14) is done in a standard way by solving $(J+3) \times (J+3)$ linear system (see [Wendland 2005], for example). In our case $J = 40$ so calculating the TPS φ is possible at interactive rates.

In the rest of this section we will describe how to set the interpolation constrains $(c_j, d_j), j = 1, .., J$ for the TPS φ so that f defined in eq.(12) will provide smooth transition to the similarities defined at its edges. Note, that although more elaborate basis functions can used to prescribe derivative information along the edges, we found that using TPS with the following discretization of the constraints to work well in practice. We discretize these constrains by spreading K (we took $K = 10$ in our experiments) equally spread points $\mathcal{P}_\alpha^Z, \mathcal{P}_\beta^Z, \mathcal{P}_\alpha^W, \mathcal{P}_\beta^W$ on each of the edges $e_\alpha^Z, e_\beta^Z, e_\alpha^W, e_\beta^W$ (resp.). To control the derivatives we also add a second line of points, called offset points, for each edge (see inset figure below, top-left). We create the offset points by creating a copy for each point on the edges and translating it a certain distance in the direction of the inward normal to the edge. Let us denote by $\mathbf{n}_\alpha^Z, \mathbf{n}_\beta^Z$ the inward normal of the edges e_α^Z, e_β^Z (resp.), and similarly for the quadrilateral W. Then for every point $p \in \mathcal{P}_\alpha^Z$ we define its offset point $\widetilde{p}$ by

$$\widetilde{p} = p + \mathbf{n}_\alpha^Z \, \delta \Big[\, |z_\alpha - z_{\beta+1}| \, |p - z_{\alpha+1}| + |z_\beta - z_{\alpha+1}| \, |p - z_\alpha| \, \Big], \tag{15}$$

where $\delta > 0$ is a parameter setting the relative distance between the two lines (in our experiments we use $\delta = 0.01$). The reason we use linear interpolation of the distances between the two edges $|z_\alpha - z_{\beta+1}|$, $|z_\beta - z_{\alpha+1}|$ is to avoid cases of conflict between the derivative constraints when the edges are transformed close to one another. In other words, we set the normal derivative to be proportional to the prescribed derivative by the edge's similarity transformation. We do the same for the rest of the constrained edges. Lastly, we move these point constraints to the suitable (Möbius) coordinate system by transforming the points via m_Z or m_W: $\{c_j\} = m_Z(\mathcal{P}_\alpha^Z \cup \mathcal{P}_\beta^Z)$ and $\{d_j\} = m_W(\mathcal{P}_\alpha^W \cup \mathcal{P}_\beta^W)$. In the inset figure we show in the bottom-right the final point constraints for one quadrilateral ($\{c_j\}$ for a source quadruplet, or $\{d_j\}$ for target quadruplet). The pseudocode for calculating the constrained FPI's components (m_Z, m_W, φ) is provided in Algorithm 3.

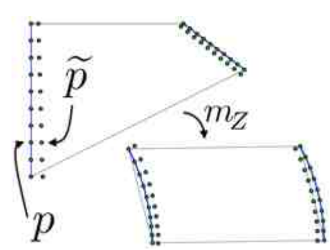

Algorithm 3: deformed_FPI(Z,W)

Input: Source points $Z = \{z_1, z_2, z_3, z_4\}$,
Target points $W = \{w_1, w_2, w_3, w_4\}$,
Constrained edges α, β,
Line offset parameter δ,
Number of offset points per edge $K = 10$
Output: deformed-FPI transformation $f(z) = m_2 \circ \varphi \circ m_1$

m_Z = quadruplet_to_parallelogram(Z)
m_W = quadruplet_to_parallelogram(W)
spread points $\mathcal{P}_\alpha^Z, \mathcal{P}_\beta^Z$, and $\mathcal{P}_\alpha^W, \mathcal{P}_\beta^W$
forall $\vartheta \in \{\alpha, \beta\}$, $\Sigma \in \{Z, W\}$ **do**
  $\mathcal{P}_\vartheta^\Sigma$ = $\mathcal{P}_\vartheta^\Sigma \cup$ offset($\mathcal{P}_\vartheta^\Sigma, Z, W, \delta$)
end
$\{c_j\}_{j=1}^K$ = $m_Z(\mathcal{P}_\alpha^Z \cup \mathcal{P}_\beta^Z)$
$\{d_j\}_{j=1}^K$ = $m_W(\mathcal{P}_\alpha^W \cup \mathcal{P}_\beta^W)$
φ = calculate_TPS_coefficients($\{c_j\}$, $\{d_j\}$)
Return $f = m_W^{-1} \circ \varphi \circ m_Z$

Note that these constraints (bottom-right in the inset figure) are close to the parallelogram's edges and are still uniformly spread after transformed by the Möbius transformation m_Z- this means that the FPI is a good approximation to the desired deformation of the ROI and that only a rather small extra deformation over the affine map is needed to adjust to the edge constraints.

Figure 7 demonstrates the use of the two types of operators to deform 2D shapes.

5. RESULTS

In this section we investigate the performance of the 4-point interpolant (basic FPI) scheme, as well as its application to compound deformations (constrained FPI), and compare to a variety of previous work.

5.1 Basic FPI deformations.

Figure 8 demonstrates several deformations of a square domain with 4 control points placed at its corners. The top row illustrates the interpolation constraints in every column, and each following row depicts the result of one particular algorithm: FPI is the 4-point interpolant introduced in this paper, BIL is bilinear interpolation, PROJ is projective transformation (both BIL, PROJ are common 4-point based warps), MVC is Mean Value Coordinates [Floater 2003; Ju et al. 2005], MLS-SIM is Moving Least-Squares deformations with similarity transformations [Schaefer et al. 2006], LSCM is Least-Squares Conformal Maps [Lévy et al. 2002; Igarashi et al. 2005], CG-P2P is Cauchy-Green coordinates with point interpolation constraints [Weber et al. 2009], and ARAP is As-Rigid-

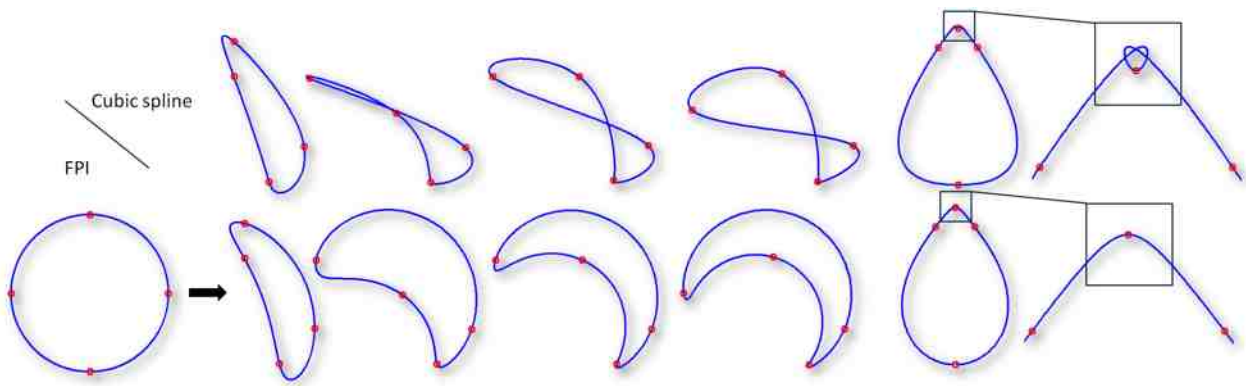

Fig. 6: *Simple shapes with FPI: we show curve fitting through four points (red). On top we use cubic spline to interpolate the four points. On bottom we use FPI to deform a perfect circle. Note that the curve fitted by FPI will never crosses itself and has intuitive behavior.*

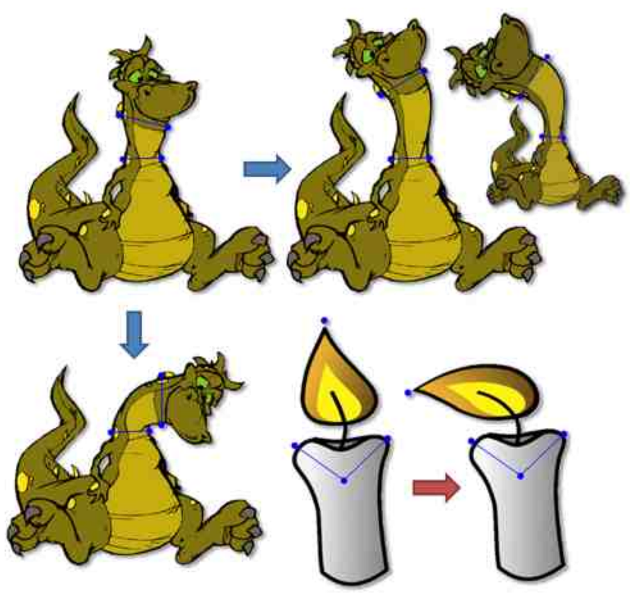

Fig. 7: *The two deformation operators: the user marks 4 points (blue dots), and chooses which lines to constrain (blue edges). Moving the free points result in the desired deformation. (see 4).*

As-Possible shape interpolation [Igarashi et al. 2005][1]. Each deformation result shows a checkerboard pattern and a conformal *dilation* map color coded, where dark blue means zero conformal dilation and dark red 0.8. Conformal dilation is defined by $d_f = \frac{|f_{\bar{z}}|}{|f_z|} = \frac{D_f - 1}{D_f + 1}$, where D_f is the conformal distortion as defined in Section 3. Basically, the conformal dilation equals zero iff f is conformal, and otherwise is a positive value measures deviation from conformality (note that for orientation reversing maps it is larger than one, otherwise, smaller than one). We also show blow-ups of certain areas to highlight distortion and fold-overs (the same region is selected throughout each column).

Note that the conformal methods, LSCM and CG-P2P, generally have the lowest conformal dilation on average, however, at certain singular points extreme distortion can occur (see for example the blow-up of the CG-P2P conformal distortion image in the left column). Furthermore, as can be seen from these examples, conformal maps that are forced to interpolate four points will tend to introduce fold-overs (around the singular points) and extreme scaling; this phenomena can be seen visually in this figure. The table below provides quantitative comparisons of the test methods for the first column; for each method we report the maximal conformal dilation, mean conformal dilation, standard-deviation of conformal dilation, and area-distortion (scale) measure which is defined by $\max J_f + 1/\min J_f$, where J_f is the jacobian of the map f.

	max d_f	mean d_f	std d_f	area dist
FPI	**0.32**	**0.32**	**0**	**5.28**
BIL	0.79	0.4	0.15	7.27
PROJ	0.81	0.67	0.11	61.63
MVC	0.71	0.41	0.14	5.26
MLS-SIM	0.82	0.33	0.18	10.10
LSCM	1.70	0.05	0.09	307.7
CG-P2P	3.04	0.04	0.12	314.2
ARAP	1.64	0.17	0.13	93.38

Note that the FPI has minimal maximum conformal dilation among all the method tested. Furthermore, it has constant conformal dilation, as the standard deviation vanishes. Note for example that LSCM and CG-P2P have lower mean conformal dilation, however their maximal conformal dilation is high; these methods are "perfectly" conformal *except* at a few singular points, the location of which is not known in advance, and in vicinity of these points the map introduces conformal distortion (e.g., vanishing complex derivative would mean that locally the map behaves like the analytic function z^n, $n \geq 2$). Furthermore, in the vicinity of these singular points the map introduces fold-overs and extreme scaling (the latter can occur at other places as-well). High area distortion means that the jacobian is unbounded or close to zero, and in the case of MLS-SIM, LSCM, CG-P2P, and ARAP implies that we are close to singularity at-least at one point, which usually means a fold-over.

[1] In our implementation we used the LSCM rotation field to seed the ARAP part, and rescaled the faces rather than rigidly fit them as it led to better results in our experiments.

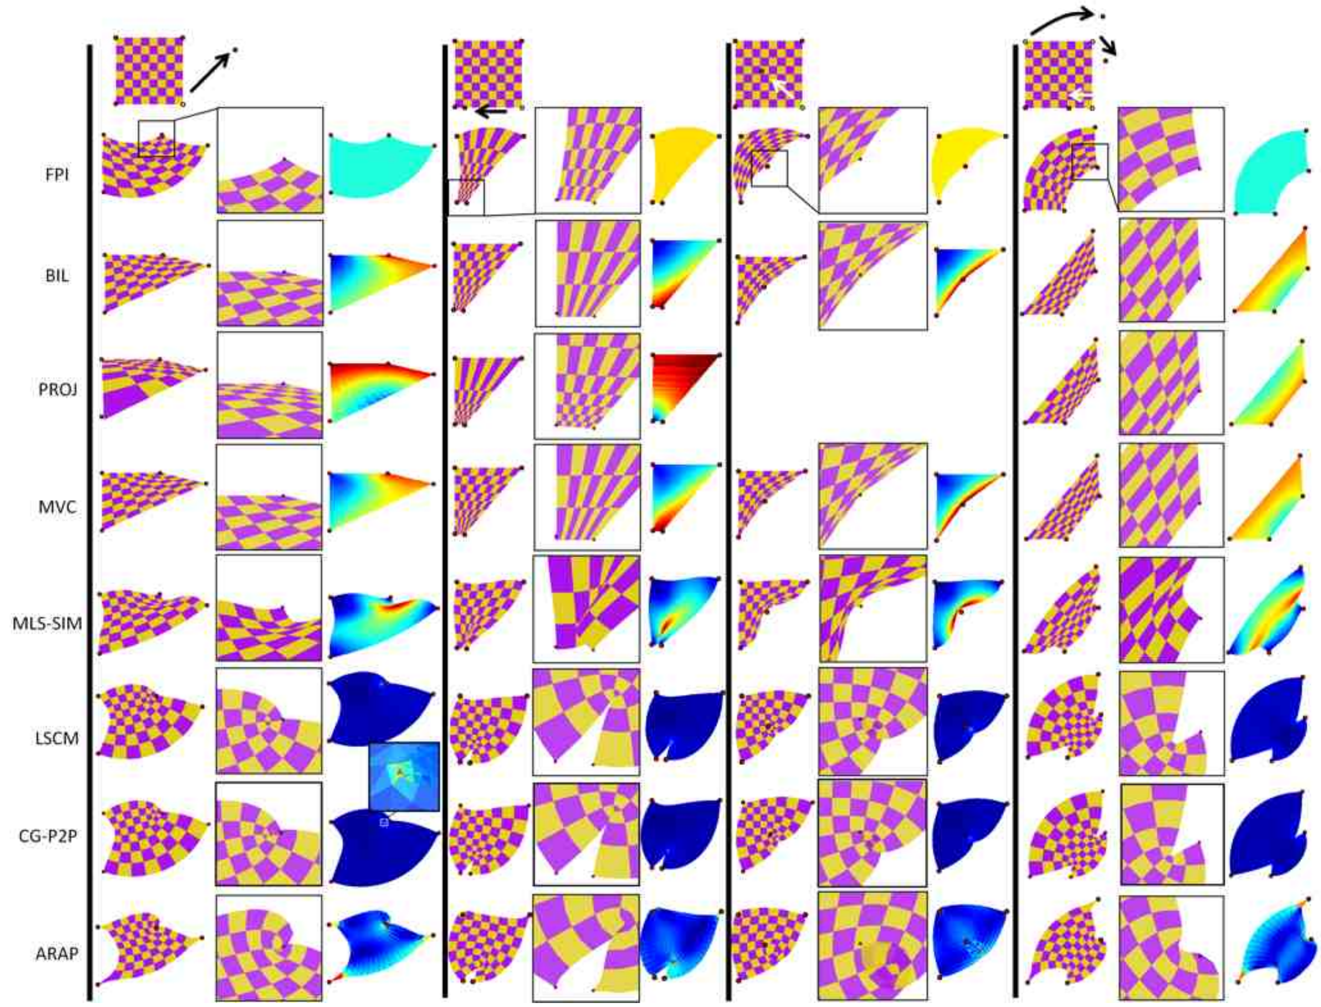

Fig. 8: *Four examples of deformations of a square domain guided with four interpolation points placed at its corners (left column). Depicted are the results of the following methods (top to bottom): 4-point Interpolant (FPI, this paper), Bilinear warping (BIL), Projective warping (PROJ), Mean Value Coordinates (MVC), Moving Least-Squares with similarities (MLS-SIM), Least-Squares Conformal Maps (LSCM), Cauchy-Green coordinates with point to point (CG-P2P), and As-Rigid-As-Possible deformation (ARAP). For each example we show a checkerboard pattern and color-coded conformal distortion (dilation) image (blue is low, red is high distortion). Note that the FPI has a constant conformal distortion, lower than the maximal conformal distortion of the other maps. Also note the insets showing areas where the distortion is high. LSCM and CG-P2P are both conformal maps (approximated for LSCM) and therefore generally have zero conformal distortion, however, they introduce extreme scaling and fold-overs in vicinity of singularities where the maps fails to be conformal (vanishing complex derivative), see for example the inset figure showing conformal distortion near such singularity.*

FPI is a simple 4-point deformation operator that can be used to intuitively manipulate simple shapes. Figure 6 (bottom), for example, shows a simple application of the FPI formula for constructing non-self-intersecting curves passing through four anchor points (red circles); the curve was created by mapping four equally spaced points on a circle to the prescribed anchors points. The figure also shows (top) comparison to cubic spline interpolation that does not guarantee such intersection-free behavior.

5.2 Compound FPI-based deformations

We have tested our FPI-based compound deformation algorithm in several scenarios. Figure 7 shows the basic operators on 2D shapes. Figure 9 demonstrates comparison with the conformal cage-based deformation example shown in the teaser image of [Weber 2010], which is not interpolatory, but avoids local fold-overs. The marked areas show regions of high distortion and/or extreme scaling caused by Weber's conformal method and alleviated by our method. Weber used a cage with 112 vertices for this example. We used five successive applications of our four point deformation operator, one of which is shown in the inset. Figure 10 shows another comparison, this time with both the methods of [Weber et al. 2009] (CG-P2P) and [Weber 2010]; note the zoom-ins of the leg and the red circle indicating some undesired scaling of parts of the frog's face which is a common artifact of conformal maps. In contrast, our operator is local and the head remains intact.

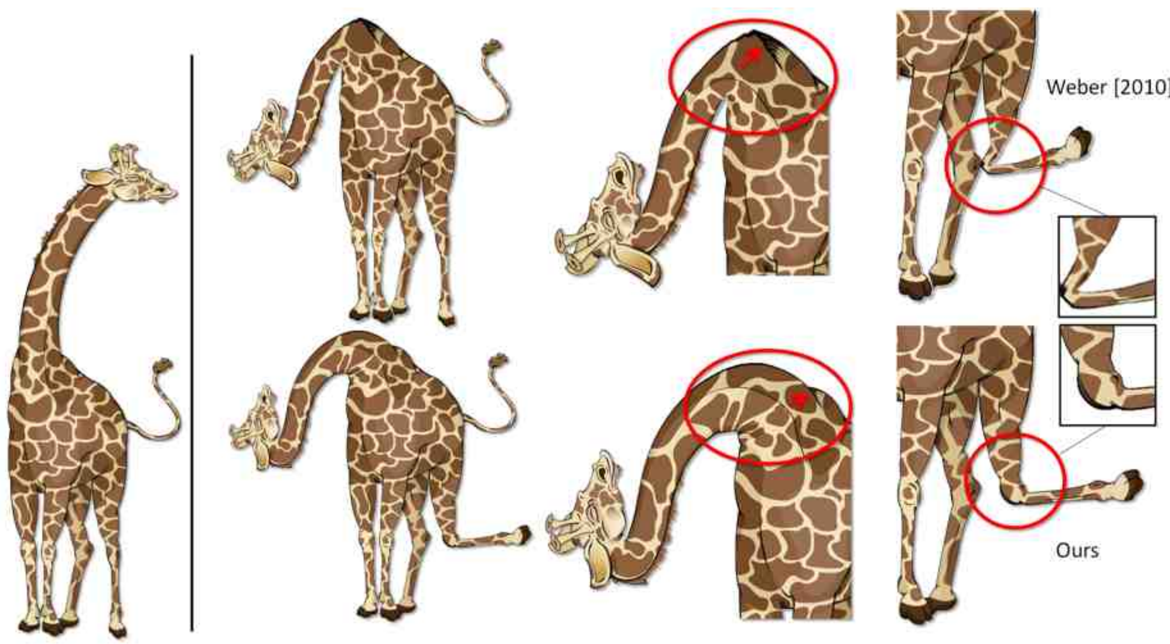

Fig. 9: *Giraffe deformation comparison: top row - the conformal algorithm of Weber and Gotsman (used as a teaser in [Weber 2010]), and bottom row - our result. In red ellipses we emphasis the main differences; the red arrow demonstrated the same area on the Giraffe's neck that was extremely scaled with the conformal map (top). The leg deformation on the right column is taken from a different image in Weber and Gotsman's teaser.*

Figure 11 compares the result of our deformed-FPI with applying the TPS interpolation directly to the positional constraints *without* performing the Möbius change of coordinates first. Note that using the TPS in the original space leads to higher conformal distortion.

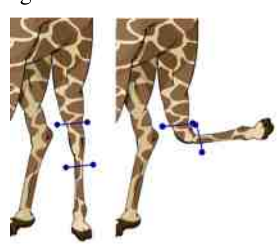

One advantage of our deformation scheme is that the user can easily and precisely control the area that will contain the conformal distortion (note that the similarities transforming the outer regions have zero conformal distortion). For comparison, in cage-based method, the distorted area is hard to control spatially, and the cage with all the needed degrees of freedom should be designed *a-priori*. Figure 12 demonstrates how choosing different ROI in deforming a human arm can create different effects: defining the ROI close to the elbow would concentrate the conformal distortion at the elbow, while taking the ROI to be the full arm will spread the conformal distortion (almost) equally across the arm. In this case the latter leads to somewhat less intuitive result as physically, the human's arm consists of rigid bones and flexible elbow.

The deformation method we suggest in this paper consists of simple closed-form formulas: eq.(1) for the basic FPI, and eq.(12) for the constrained FPI. The coefficients in these formulas are computed via algorithms 2 and 3 (resp.), and then each point is deformed by the corresponding analytical formula. The algorithm is extremely simple to implement. Furthermore, it is efficient computationally; in this paper we computed the deformations on a triangulation of the domain and texture mapped the images. All our meshes used in this paper consisted of maximal number of 10K vertices. The basic FPI scheme takes 0.001s to deform 1K vertices on 2.4 GHz processor. The constrained FPI requires additional TPS computation over the basic FPI. In our implementation we used 40 centers for the TPS and were required to solve 40×40 linear system; for 1K vertices computing the deformation and applying it takes 0.0016s on the same processor. Note that the overhead of the TPS is minor.

6. DISCUSSION, LIMITATIONS AND FUTURE WORK

We have presented a simple formula for 4-point planar warping that spreads conformal distortion equally and has optimal worst-case conformal distortion properties.

We have shown that the FPI can be used for building deformation operators that are simple and can provide an alternative to previous planar warping and interpolation methods. In particular the benefits over the more common cage-based techniques are: 1) the user can define the deformed region on-the-fly, and does not need to design an entire cage with enough degrees of freedom in a separate preprocess stage, 2) the mapping comes with certain guarantees, 3) the algorithm is very simple, consisting of a formula that describes the mapping, 3) the deformation is local - the user control precisely the area to be deformed (this requirement is often raised by end-users), and 4) the FPI has 4 control points which we found very intuitive to define deformations.

The method described in the paper has some limitations. First, in our current implementation, the constrained deformation (Section 4) is described only for ROIs bounded by straight lines. However, generalizing this operator to consider ROIs bounded by *any* curve connecting adjecent control points is trivial - there is nothing in our construction that builds on the fact that the constrained edges

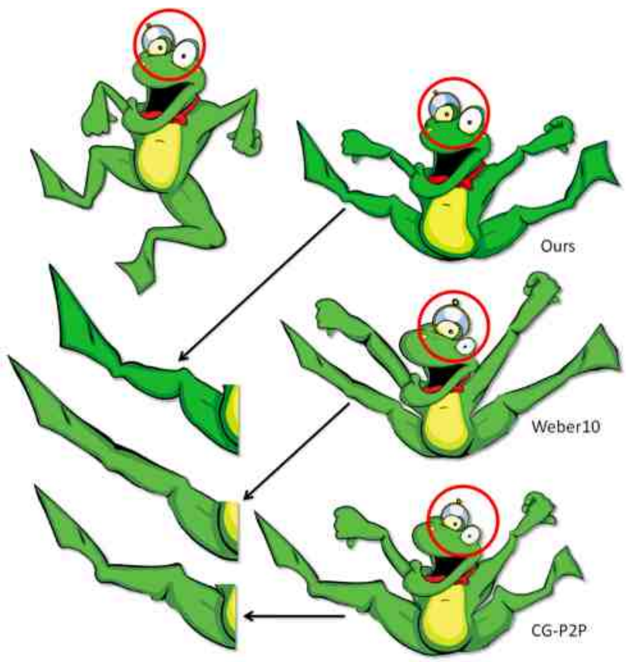

Fig. 10: *Articulation of a frog. We compare to Weber10 [2010], and CG-P2P [Weber et al. 2009]. Note the zoom-ins of the frog's right leg, and the red circle indicating undesired scaling in the two bottom examples.*

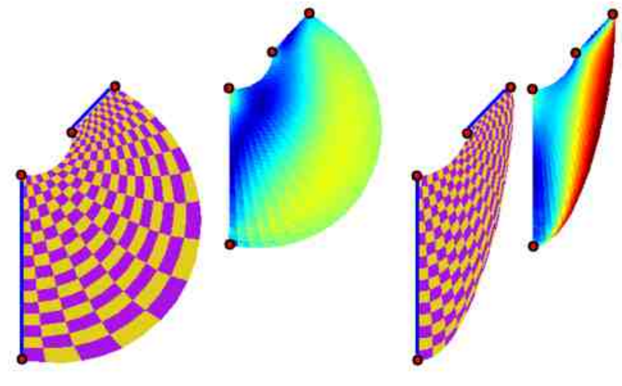

(a)

(b)

Fig. 11: *Comparison of deformation in a square ROI using the constrained FPI (left column) versus TPS with the same edge constraints – eq.(15) (right column).*

are straight. The second limitation is that the contrained deformation does not allow simultaneous control over adjacent edges of the ROI; the similarities defined on different edges of the ROI will not match in general using our model, and therefore a more complicated model should be used to constrain the deformation outside the ROI when edges being manipulated by the user share a point.

As for future work, we would be interested to find optimal quasi-conformal mapping in different spaces than the periodic mappings. One interesting example is to consider the collection of maps between the straight-edged quadrilateral that interpolate the corners. Another example is the sphere. Also, finding provably optimal quasi-conformal mapping with derivative constraints would be interesting for our application; currently, we are using the FPI as our approximation for such optimal map. Lastly, we would like to develop a 4-point deformation application for touch-screens (currently we have a standard PC implementation) as we believe that humans will find 4-points based deformation intuitive and useful (using two finger out of each hand).

Acknowledgements

We would like to thank Ofir Weber for supplying the images of Figures 9,10, and Mirela Ben-Chen for the CG-P2P result in Figure 8. Also, we thank Google, Adobe, Intel, NSERC and NSF (CCF-0937139).

REFERENCES

AHLFORS, L. 1966. *COMPLEX ANALYSIS*.

AHLFORS, L. V. 2006. *Lectures on Quasiconformal Mappings*. University Lecture Series, vol. 38.

BOOKSTEIN, F. L. 1989. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell. 11*, 567–585.

DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. Intrinsic Parameterizations of Surface Meshes. *Computer Graphics Forum 21*.

FLETCHER, A. AND MARKOVIĆ, V. 2007. *Quasiconformal maps and Teichmüller theory*. Oxford graduate texts in mathematics. Oxford University Press.

FLOATER, M. S. 2003. Mean value coordinates. *Comput. Aided Geom. Des. 20*, 19–27.

IGARASHI, T., MOSCOVICH, T., AND HUGHES, F. J. 2005. As-rigid-as-possible shape manipulation. *ACM Trans. Graph 24*, 1134–1141.

JU, T., SCHAEFER, S., AND WARREN, J. 2005. Mean value coordinates for closed triangular meshes. *ACM Trans. Graph. 24*, 561–566.

LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLO T, J. 2002. Least squares conformal maps for automatic texture atlas generation. In *ACM SIGGRAPH conference proceedings*, ACM, Ed.

LIPMAN, Y., LEVIN, D., AND COHEN-OR, D. 2008. Green coordinates. *ACM Trans. Graph. 27,* 3.

SCHAEFER, S., MCPHAIL, T., AND WARREN, J. 2006. Image deformation using moving least squares. *ACM Trans. Graph. 25*, 533–540.

SEDERBERG, T. AND PARRY, S. 1986. Free-form deformation of solid geometric models. *SIGGRAPH Comput. Graph. 20*, 151–160.

WEBER, O., G. C. 2010. Controllable conformal maps for shape deformation and interpolation. *ACM Trans. Graph. 29*, 78:1–78:11.

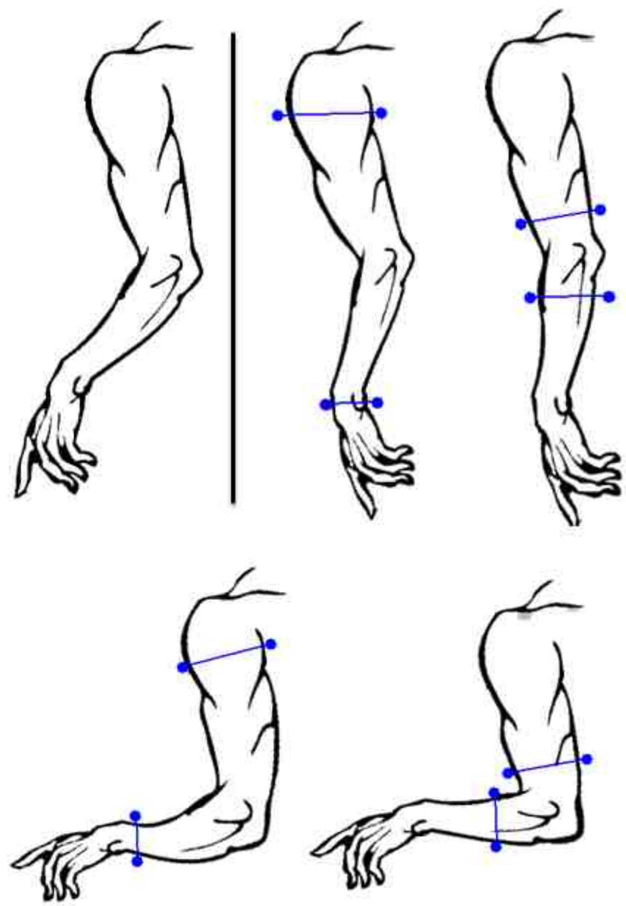

Fig. 12: *Controlling the locality of the deformation. Using different choices of ROI the user can presciently determine where the conformal distortion in the deformation will be concentrated. We show an image of a human's hand (top-left), and two deformations (top-right, and bottom row). Note that prescribing the FPI edges near the elbow concentrate the deformation's distortion in that area while taking edges farther from the elbow causes the conformal distortion to distribute along the entire arm.*

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

WENDLAND, H. 2005. *Scattered Data Approximation.* Cambridge Monographs on Applied and Computational Mathematics (No. 17).

ZENG, W. AND GU, X. D. 2011. Registration for 3d surfaces with large deformations using quasi-conformal curvature flow. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on.* 2457 –2464.

ZENG, W., LUO, F., YAU, S. T., AND GU, X. D. 2009. Surface quasi-conformal mapping by solving beltrami equations. In *Proceedings of the 13th IMA International Conference on Mathematics of Surfaces XIII.* Springer-Verlag, Berlin, Heidelberg, 391–408.

ZENG, W., MARINO, J., CHAITANYA GURIJALA, K., GU, X., AND KAUFMAN, A. 2010. Supine and prone colon registration using quasi-conformal mapping. *Visualization and Computer Graphics, IEEE Transactions on 16,* 6 (nov.-dec.), 1348 –1357.

APPENDIX

Appendix A.

In this appendix we provide the proof that the periodic map $f \in \mathcal{F}$ that bijectively takes one parallelogram $P(\eta, \xi)$ to another $P(\widehat{\eta}, \widehat{\xi})$ with lowest maximal conformal distortion is the affine map. This fact, although seems natural, is not trivial to prove. The proof is contained within Ahlfors [Ahlfors 2006] proof of a slightly different problem. We decided to adapt the proof to our setting for two reasons: first, it has ideas that we believe can stimulate researchers to think about the type of problem discussed in this paper in a more general context, and second, the ideas are folded inside Ahlfors arguments and are not easily accessible.

Since we can use the conformal map $z \mapsto z/\xi$ to map (without introducing conformal distortion) $P(\eta, \xi)$ to $P(\tau = \frac{\eta}{\xi}, 1)$, we will only consider parallelograms of the form $P(\tau) := P(\tau, 1)$. W.l.o.g we can assume $\mathrm{Im}\,(\tau) > 0$. Given a differentiable map between two periodic parallelograms (that interpolates the corners) $f : P(\tau) \to P(\widehat{\tau})$ we will measure its maximal conformal distortion by $K_f = \max_{z \in P(\tau)} D_f(z)$. We show that the map $f \in \mathcal{F}$ that minimizes K_f is the affine map taking $\tau \to \widehat{\tau}$ and fixing 1.

In Lemma A.1, we prove that any differentiable map $f : P(\tau) \to P(\widehat{\tau})$ must satisfy

$$K_f \geq \frac{\mathrm{Im}\,(\tau)}{\mathrm{Im}\,(\widehat{\tau})}. \tag{16}$$

Given this lemma we will show the result. We note that given any $a, b, c, d \in \mathbb{Z}$ such that $\det \begin{pmatrix} a & b \\ c & d \end{pmatrix} = 1$ (unimodular matrix) the periodic parallelogram $P(a\tau + b, c\tau + d)$ is exactly equivalent to $P(\tau) = P(\tau, 1)$. It can be thought of as different parametrization of the same object. Similarly, $P(a\widehat{\tau} + b, c\widehat{\tau} + d)$ is equivalent to $P(\widehat{\tau})$. Note that $f(a\tau + b) = a\widehat{\tau} + b$, and $f(c\tau + d) = c\widehat{\tau} + d$, and in general f satisfies $f\left(z + k(a\tau + b) + \ell(c\tau + d)\right) = f(z) + k(a\widehat{\tau} + b) + \ell(c\widehat{\tau} + d)$.

Next, let us apply the (similarity) transform $S_1 : z \mapsto z/(c\tau + d)$ to map $P(a\tau + b, c\tau + d)$ to $P(\frac{a\tau+b}{c\tau+d})$, and $S_2 : z \mapsto z/(c\widehat{\tau} + d)$ to map $P(a\widehat{\tau} + b, c\widehat{\tau} + d)$ to $P(\frac{a\widehat{\tau}+b}{c\widehat{\tau}+d})$. The map $\widetilde{f} = S_2 \circ f \circ S_1^{-1}$ maps $P(\frac{a\tau+b}{c\tau+d})$ to $P(\frac{a\widehat{\tau}+b}{c\widehat{\tau}+d})$, and satisfies $\widetilde{f}(z + k\frac{a\tau+b}{c\tau+d} + \ell) = \widetilde{f}(z) + k\frac{a\widehat{\tau}+b}{c\widehat{\tau}+d} + \ell$. Furthermore, $D_{S_2 \circ f \circ S_1^{-1}}(z) = D_f(S_1^{-1}(z))$ for all $z \in \mathbb{C}$. Applying the lower bound (16) then implies

$$K_f = K_{S_2 \circ f \circ S_1^{-1}} \geq \mathrm{Im}\left(\frac{a\tau + b}{c\tau + d}\right) \Big/ \mathrm{Im}\left(\frac{a\widehat{\tau} + b}{c\widehat{\tau} + d}\right), \tag{17}$$

for all $a, b, c, d \in \mathbb{Z}$ s.t. $ad - bc = 1$. To finish this argument, Ahlfors uses the following elegant geometrical observation: let m be a Möbius transformation taking the upper-half plane to the interior of the unit disk such that $m(\widehat{\tau}) = 0$. Denote $h(z) = \frac{az+b}{cz+d}$, where $a, b, c, d \in \mathbb{Z}$, $ad - bc = 1$. From (17) we know that

$$\mathrm{Im}\,(h(\widehat{\tau}))\, K_f \geq \mathrm{Im}\,(h(\tau))\,.$$

Denote the set $C = \{z \mid \mathrm{Im}\,(z) > \mathrm{Im}\,(h(\widehat{\tau}))\, K_f\}$. The bound above implies that $m^{-1}(\tau)$ is not inside the open circle $m^{-1} \circ h^{-1}(C)$. Furthermore, the shortest hyperbolic distance between $h(\widehat{\tau})$ and the closure of C is $d_H(h(\widehat{\tau}), C) = d_H\left(\mathrm{i}\mathrm{Im}\,(h(\widehat{\tau}))\,, \mathrm{i}K_f \mathrm{Im}\,(h(\widehat{\tau}))\right) = \log(K_f)$. This is also the hyperbolic distance between the origin and the circle $m^{-1} \circ h^{-1}(C)$ in the hyperbolic disk. The circle $m^{-1} \circ h^{-1}(C)$ is osculating to the boundary of the unit disk at the point $m^{-1} \circ h^{-1}(\infty)$. Since we can always find unimodular h such that $h^{-1}(\infty)$ is an arbitrary rational number, the circle C can osculate to a dense set of points on the boundary of the unit disk. Since $m^{-1}(\tau)$ cannot be inside any of these circles the hyperbolic distance of $m^{-1}(\tau)$ to the origin should be less or equal to the distance of any such circle to the origin which we already computed to be $\log(K_f)$. We conclude that

$$d_H(\tau, \widehat{\tau}) \leq \log(K_f).$$

Let us show that the affine map $f : P(\tau) \to P(\widehat{\tau})$ defined by

$$f(z) = \frac{(\widehat{\tau} - \overline{\tau})\, z + (\tau - \widehat{\tau})\, \overline{z}}{\tau - \overline{\tau}}$$

has $K_f = e^{d_H(\tau,\widehat{\tau})}$. Indeed,

$$K_f = D_f = \frac{|f_z| + |f_{\overline{z}}|}{|f_z| - |f_{\overline{z}}|} = \frac{|\widehat{\tau} - \overline{\tau}| + |\widehat{\tau} - \tau|}{|\widehat{\tau} - \overline{\tau}| - |\widehat{\tau} - \tau|} = e^{d_H(\widehat{\tau},\tau)}.$$

■

LEMMA A.1. *Let $f : P(\tau) \to P(\widehat{\tau})$ be a differentiable map. Then,*

$$K_f \geq \frac{\mathrm{Im}\,(\tau)}{\mathrm{Im}\,(\widehat{\tau})}.$$

Although is possible to prove this lower bound with extremal length method, we will use a more direct technique due to Grötzch. Given a parallelogram $P(\tau)$ we parameterize it over the unit square by $z = s\tau + t$, $0 \leq s, t \leq 1$. Then the change of variable formula implies

$$\iint_{P(\tau)} \phi(z) dx\, dy = \mathrm{Im}\,(\tau) \int_0^1 \int_0^1 \phi(s\tau + t) ds\, ts, \qquad (18)$$

for any integrable ϕ. Next, fix s and consider the curve $\gamma_s(t) = s\tau + t$, $0 \leq t \leq 1$. We have

$$1 \leq \mathrm{length}(f(\gamma_s)) = \int_0^1 |f_z(\gamma_s(t)) + f_{\overline{z}}(\gamma_s(t)|\, dt \leq \int_0^1 |f_z| + |f_{\overline{z}}|\, dt.$$

Integrating both sides w.r.t to $s \in [0, 1]$, multiplying both sides by $\mathrm{Im}\,(\tau)$ and using (18) we get $\mathrm{Im}\,(\tau) \leq$

$$\iint_{P(\tau)} |f_z| + |f_{\overline{z}}|\, dx\, dy = \iint_{P(\tau)} \frac{|f_z| + |f_{\overline{z}}|}{\sqrt{|f_z|^2 - |f_{\overline{z}}|^2}} \sqrt{J_f} dx\, dy,$$

where in the last equality we multiplied and divided by the square-root of the jacobian J_f of f. Using Cauchy-Schwarz:

$$\iint_{P(\tau)} \frac{|f_z| + |f_{\overline{z}}|}{\sqrt{|f_z|^2 - |f_{\overline{z}}|^2}} \sqrt{J_f} \leq \left[\iint_{P(\tau)} \frac{|f_z| + |f_{\overline{z}}|}{|f_z| - |f_{\overline{z}}|}\right]^{\frac{1}{2}} \left[\iint_{P(\tau)} J_f\right]^{\frac{1}{2}},$$

squaring both sides and using previous inequality we get,

$$\mathrm{Im}\,(\tau)^2 \leq \left[\iint_{P(\tau)} D_f\right] \mathrm{Im}\,(\widehat{\tau}) \leq K_f \mathrm{Im}\,(\tau)\, \mathrm{Im}\,(\widehat{\tau})$$

rearranging the terms proves the lemma.■