

Hyperbolic Orbifold Tutte Embeddings

Noam Aigerman Yaron Lipman
Weizmann Institute of Science

Abstract

Tutte’s embedding is one of the most popular approaches for computing parameterizations of surface meshes in computer graphics and geometry processing. Its popularity can be attributed to its simplicity, the guaranteed bijectivity of the embedding, and its relation to continuous harmonic mappings.

In this work we extend Tutte’s embedding into hyperbolic cone-surfaces called orbifolds. Hyperbolic orbifolds are simple surfaces exhibiting different topologies and cone singularities and therefore provide a flexible and useful family of target domains. The hyperbolic Orbifold Tutte embedding is defined as a critical point of a Dirichlet energy with special boundary constraints and is proved to be bijective, while also satisfying a set of points-constraints. An efficient algorithm for computing these embeddings is developed.

We demonstrate a powerful application of the hyperbolic Tutte embedding for computing a consistent set of bijective, seamless maps between all pairs in a collection of shapes, interpolating a set of user-prescribed landmarks, in a fast and robust manner.

Keywords: Tutte embedding, hyperbolic, orbifold, discrete harmonic, injective parameterization, surface mapping

Concepts: •Computing methodologies → Mesh models;

1 Introduction

Surface parameterization – the task of mapping a surface-mesh to a simpler two-dimensional domain – is a central tool in Geometry Processing, Computer Graphics and Scientific Computation.

One of the most fundamental parameterization algorithms is Tutte’s embedding of planar graphs [Tutte 1963]. Tutte’s embedding has been generalized to convex-combination maps [Floater 2003a] and has been used extensively for parameterizing polygonal meshes [Desbrun et al. 2002; Floater 2003b; Gortler et al. 2006; Weber and Zorin 2014]. The key benefit in this method is that the resulting map is guaranteed to be bijective, and is a discrete version of a harmonic map, inheriting powerful properties from the smooth case, such as the maximum principle.

So far, Tutte’s embedding has been used for mapping surfaces into flat, *i.e.*, Euclidean, target domains: Convex polygonal regions in the Euclidean plane [Tutte 1963; Floater 2003a], the flat torus [Lovász 2004; Gortler et al. 2006], and Euclidean orbifolds, which are flat surfaces with cone-singularities, including spheres with 3-4 cones [Aigerman and Lipman 2015]. Currently, there are no gener-

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. © 2016 ACM.

SA ’16 Technical Papers, December 05-08, 2016, , Macao

ISBN: 978-1-4503-4514-9/16/12

DOI: <http://dx.doi.org/10.1145/2980179.2982412>




Figure 1: The Tutte embedding of a mesh into a hyperbolic orbifold. The embedding into the basic tile (highlighted in bold color) can tile the entire Poincaré disk (tiling shown in faded colors), exhibiting the orbifold’s special symmetry properties.

alizations of Tutte’s embedding into non-Euclidean target domains and/or domains with more than 4 cones.

The goal of this paper is to generalize Tutte’s embedding to *hyperbolic orbifolds* – a family of hyperbolic surfaces (*i.e.*, 2-manifolds with constant negative curvature) possessing an arbitrary number of cones. Formally, a hyperbolic orbifold is defined as the quotient of the hyperbolic plane with respect to a hyperbolic symmetry group. There is an infinite number of hyperbolic orbifolds, exhibiting an extremely wide variety of cone structures and topologies, making them good target domains for surface parameterizations. We shall denote a hyperbolic orbifold as $\mathcal{O}$.

Our approach to defining and computing Hyperbolic Orbifold Tutte embeddings is simple and inspired by the Euclidean case [Aigerman and Lipman 2015]: Look for a convex-combination map¹ in the hyperbolic plane satisfying boundary conditions that ensure the hyperbolic plane can be perfectly tiled using the target domain; this ensures the image of the mapping is the desired hyperbolic orbifold. We formulate this problem as a smooth optimization problem, consisting of minimizing the Dirichlet energy while adhering to the orbifold’s boundary constraints. We prove that a critical point of this problem is a bijective map, mapping user-prescribed vertices to the cones of the target orbifold. To compute this critical point we adapt a quasi-Newton method. Figure 1 demonstrates a hyperbolic Tutte embedding of a sphere-type mesh of an owl into a 7-cone sphere-type orbifold; the colored spheres indicate the points on the model mapped to the orbifold’s cones (colored with matching colors).

As a main practical application, we demonstrate the usefulness of the Hyperbolic Orbifold Tutte embeddings for *surface-to-surface* mappings. This is achieved by embedding the surfaces into the same target hyperbolic orbifold. The resulting map is guaranteed to

¹A convex combination map places each vertex in a weighted average of its 1-ring neighbors (see [Floater 2003a]).




Figure 2: The hyperbolic orbifold Tutte embedding algorithm enables computing a collective homeomorphic mapping between one hundred models from the FAUST [Bogo et al. 2014] data-set. Each individual map was computed in 80 seconds in a fully-parallelizable process.

be bijective, seamless, and to interpolate any set of prescribed landmarks. Figure 2 depicts a collective mapping of the entire FAUST dataset [Bogo et al. 2014], computed using our algorithm; the landmarks are visualized as colored spheres. The color and texture of the meshes are transferred from one of the models (fourth person from the right, first row) to the rest of the set. Figure 3 shows a single pair from this collection in more detail. As all of the mappings are defined via a single common domain, the mappings are *cycle-consistent* [Nguyen et al. 2011], by construction. Cycle consistency is a necessary and sufficient condition for well-defining correspondences of points across a collection of surfaces, *i.e.*, the meshes can be consistently colored, textured or labeled.

Tsui *et al.* [2013] were the first to advocate the useful idea of using hyperbolic orbifolds as target domains to facilitate surface-to-surface mapping: they embed each surface into a conformally-equivalent hyperbolic orbifold using the CETM parameterization algorithm [Springborn et al. 2008] and map the two orbifolds using a harmonic map. Although powerful, CETM may fail to produce a valid embedding, in turn forcing the approach of Tsui *et al.* to fail. The observation of our work is that bijective harmonic embeddings to hyperbolic orbifolds can be computed using a Tutte-like algorithm with closed-form boundary conditions extracted directly from the hyperbolic symmetry groups.

2 Previous work

Tutte’s embedding Tutte’s embedding, [Tutte 1963; Floater 2003a], yields globally injective mappings by computing a convex-combination map into a convex polygon. [Lovász 2004; Gortler et al. 2006] extend Tutte’s embedding to the flat-torus case by integrating harmonic one-forms on the torus, and [Aigerman and Lipman 2015] extend Tutte to all Euclidean orbifold structures, *i.e.*, cone manifolds which can tile the plane, using convex combination maps with explicit boundary conditions. There are 17 Euclidean orbifolds (including the torus), with a maximum of 4 cones. In this paper we extend Tutte’s embedding further to the hyperbolic plane and thus enable embedding into (hyperbolic) orbifolds, possessing an unlimited number of cones and different topologies.

Surface parameterization Parameterizations lie at the core of Geometry Processing and Computer Graphics, for surveys see [Sheffer et al. 2006; Hormann et al. 2007]. In this context, injectiv-

ity is a sought-after property, and many papers strive to guarantee it, either locally [Hormann and Greiner 2000; Sheffer et al. 2005; Schüller et al. 2013; Aigerman et al. 2014; Weber and Zorin 2014], or globally [Lipman 2012; Campen et al. 2015; Smith and Schaefer 2015]. A prominent algorithm for computing parameterizations into the Euclidean and hyperbolic plane is conformal equivalence of triangle meshes (CETM) presented in [Springborn et al. 2008] which modifies the edge-lengths to new, conformally-equivalent, edge-lengths. However, in contrast to our method, CETM is not guaranteed to converge to a locally-injective map without incorporating edge flips and mesh-refinement, nor can it provide control over the exact target locations of the cone singularities, required for it to be incorporated in a surface-to-surface mapping application. Aside from CETM there exist other powerful algorithms for computing hyperbolic conformal parameterizations, including circle-packing [Stephenson 2005], circle patterns [Kharevych et al. 2006], and Ricci flow [Jin et al. 2008], and Euclidean cone manifold parameterization, mainly in the context of quadrangulation [Bommes et al. 2009; Myles and Zorin 2013].




Figure 3: A single map from the collection shown in Figure 2. While the maps are defined over the entire collection, each map is guaranteed to be a homeomorphism and exhibits high accuracy.




Figure 4: The two orbifold structures we use in this paper, visualized using the basic tile (and its tiling, in light color) in the Poincaré model. Left: the basic tile of a hyperbolic spherical orbifold with seven cones of π . Boundaries with the same color are associated to one-another via a Möbius transformation m ; colored points mark the cones, with cones sharing the same color associated to one another. Right: the basic tile of a hyperbolic disk orbifold with five cones of $\pi/2$. Note this pentagon has an angle of exactly $\pi/2$ at each of its five vertices. Each boundary arc is associated to itself via an anti-Möbius (reflection) transformation m . Both tiles are in fact convex in the hyperbolic plane.

Surface mappings Computing surface maps via a common base-domain is a popular approach; [Lee et al. 1999; Schreiner et al. 2004; Kraevoy and Sheffer 2004; Bradley et al. 2008] use coarse meshes as the base domain. [Aigerman et al. 2014; Weber and Zorin 2014] use a common planar domain to map disks to one another, but only require the embedding to be locally-injective. [Aigerman et al. 2015] use a flat cone manifold whose cone angles and positions change during the optimization. [Shi et al. 2013] computes a map between surfaces similarly to [Tsui et al. 2013] with the difference of using Ricci flow instead of CETM for generating a conformal parameterization, and using a pants-decomposition instead of an orbifold as a basic domain. They map domains in the hyperbolic plane using a discrete hyperbolic Dirichlet energy whose minimizers are also a hyperbolic equivalent of a convex-combination map; we use the same energy in this paper.

3 Background

We will use hyperbolic orbifolds as our target domains into which the meshes will be embedded. In this section we overview some relevant facts regarding hyperbolic geometry and define the hyperbolic orbifolds; for a more elaborate survey on hyperbolic geometry see [Cannon et al. 1997].

The hyperbolic plane $\mathbb{H}^2$ is an infinite two-dimensional surface with constant negative curvature, that is, it can be imagined as a curved plane, in which each point is a saddle point. It cannot be embedded in its entirety in $\mathbb{R}^3$ without distorting its inner distances and for that reason it is useful to work with one of its non-isometric models. A popular such model is the Poincaré disk; it models $\mathbb{H}^2$ as the open unit-disk, $D = \{z \in \mathbb{C} \mid |z| < 1\}$. Since it is not an isometric model it comes with a prescription of a non-Euclidean metric at each point $z \in D$, defined using the formula:

$$ds^2 = \frac{4}{(1 - |z|^2)^2} |dz|^2. \quad (1)$$

Since $|dz|^2$ represents the Euclidean metric, this formula can be

interpreted as follows: at $z \in D$ the angle between two tangent vectors $\xi, \eta \in \mathbb{C}$ coincides with the Euclidean angle; however, the length of a tangent vector $\xi \in \mathbb{C}$ is $4(1 - |z|^2)^{-2} |\xi|$, where $|\xi|$ is the Euclidean length of the vector ξ ; namely, the closer one is to the boundary of D the larger the scaling is.

The geodesics in the Poincaré model are circular arcs, intersecting the boundary of the disk at right angles; isometries (mappings of the hyperbolic plane preserving all intrinsic angles and distances, equivalently to rigid motions in Euclidean geometry) are the subgroup of the Möbius and anti-Möbius transformations mapping the unit disk to itself; for their explicit formula see Appendix A.

The main motivation for exploring the hyperbolic plane is that it enables geometric constructions which are, in a sense, more flexible - and thus have greater variety - than, e.g., the Euclidean plane. For instance, a fact we shall use in the construction of the boundary conditions of the orbifolds, is that for *any* triplet of angles α, β, γ whose sum is smaller than π , there exists a (unique, up to isometry) hyperbolic triangle realizing those angles.

A hyperbolic orbifold is a hyperbolic surface with cone singularities. It can be visualized as a piece of the hyperbolic plane stitched along its boundaries. For example, Figure 4, left, shows such a piece with matching boundary curves colored identically. As can be observed (by imagining stitching these boundaries physically) the resulting surface is of spherical topology, and it is hyperbolic (*i.e.*, with constant negative curvature) everywhere, except at 7 points indicated with colored disks; these points are *cone singularities*, points around which the angle sum is different than 2π - in this case the cone angles are all of π . Note that disks with the same color represent the same cone singularity. The two-dimensional layout shown in Figure 4, left, is called the *basic tile* of the orbifold, and the cone singularities are denoted as c_j . The inset depicts a conceptualization of the orbifold from Figure 4, left, embedded in three-dimensional Euclidean space; the cone singularities are indicated by colored spheres; the coloring of both the cones and the cut-curves matches the basic tile in Figure 4, left.




Figure 4 reveals the characteristic property of the hyperbolic orbifold (distinguishing it from a general hyperbolic surface with cones): it can tile the hyperbolic plane, by applying a discrete subgroup of isometric transformations (Möbius and anti-Möbius) to the basic tile; we denote such transformations as m and call this subgroup $\{m\}$ the *orbifold symmetry group*. Each m maps an arc of the basic tile to the arc it is associated with, as illustrated in Figure 4, left, for the spherical orbifold case, where each m is a Möbius transformation. In Figure 4, right, we show the basic tile of a *disk* orbifold. Here, every arc is associated to *itself*, via an anti-Möbius transformation m , reflecting the orbifold over the arc.

Formally, a hyperbolic orbifold is the quotient of the hyperbolic plane $\mathbb{H}^2$ under a discrete subgroup of the isometries of $\mathbb{H}^2$.

There is an infinite variety of hyperbolic orbifolds, exhibiting a wide range of cone structures and topologies. While the presented theory and algorithm are applicable to surfaces of all topologies, in this work we focus on embedding surfaces of two possible topologies: sphere-type, and disk-type. We will therefore use two hyperbolic orbifold sub-families: 1) sphere-type orbifolds with k cones of angle π - shown in Figure 4, left, for $k = 7$; and 2) disk-type orbifolds with k boundary cones of angle $\pi/2$ - shown in Figure 4, right, for $k = 5$. These families should satisfy the necessary condition ([Conway et al. 2008]):

$$1 - \frac{k}{4} < 0.$$

One simple conclusion from this equation is that any hyperbolic orbifold of these two types has at-least 5 cones. It is quite remarkable that the above condition is also sufficient: for each k satisfying this equation there exists a corresponding hyperbolic orbifold.

4 Hyperbolic Orbifold Tutte embedding

We assume to be given a 3-connected triangular mesh $\mathbf{M}$, and a set of k vertex landmarks $p_1, p_2, \dots, p_k$. Our goal is to compute an embedding of $\mathbf{M}$ into a hyperbolic orbifold $\mathcal{O}$ with a matching number of cones, so that each landmark will be mapped to a corresponding cone. The embedding is defined using a simple principle, generalizing the Euclidean case in [Aigerman and Lipman 2015]:

1. If the mesh is not a disk, cut it open to a disk mesh, $\mathbf{M}^\circ$.
2. Build a convex combination map into the hyperbolic plane, $\Phi : \mathbf{M}^\circ \rightarrow \mathbb{H}^2$ under appropriate boundary conditions that ensure the image of the map is the desired hyperbolic orbifold.

We now elaborate on these two steps.

Cutting In case $\mathbf{M}$ is not a disk (*i.e.*, a sphere in all examples in this paper), we cut it sequentially through the landmarks, $p_1 \rightarrow p_2 \rightarrow \dots \rightarrow p_k$ to a disk mesh, $\mathbf{M}^\circ = (\mathbf{V}, \mathbf{E}, \mathbf{T})$ with vertices $\mathbf{V}$, edges $\mathbf{E}$ and triangles $\mathbf{T}$. As we cut the mesh, each vertex v_i along the cut (except for the first and last cones, p_1, p_k) is duplicated to two vertices $v_i, v_{i'}$ lying on opposite sides of the boundary of the disk mesh. A landmark and its duplicate are similarly denoted $p_j, p_{j'}$.

In case the input is already a disk-type mesh we denote $\mathbf{M}^\circ = \mathbf{M}$ and denote $v_i = v_{i'}$ for boundary, non-landmark vertices.

Embedding The disk mesh $\mathbf{M}^\circ$ is embedded into the hyperbolic plane by mapping its vertices into the Poincaré disk via a map $\Phi : \mathbf{V} \rightarrow D$. The image of each vertex is represented by an unknown complex number, denoted $\Phi_i = \Phi(v_i) \in D$. This map can be then extended to the edges $e_{ij} \in \mathbf{E}$ by drawing geodesic (circular) arcs between Φ_i, Φ_j .

We now replace the Euclidean convex-combination maps used in the Euclidean case with *hyperbolic* convex-combination maps. Convex-combinations in the hyperbolic plane can be defined via the Karcher mean [Karcher 1977]. This naturally leads to defining the convex-combination map as a critical point of the discrete Dirichlet energy [Shi et al. 2013; Tsui et al. 2013],

$$E(\Phi) = \frac{1}{2} \sum_{(i,j) \in \mathbf{E}} w_{ij} d(\Phi_i, \Phi_j)^2, \quad (2)$$

where $d(z, w)$ is the hyperbolic distance between two points $z, w \in D$, and w_{ij} are positive weights assigned to the edges (see appendix A for the formula of the hyperbolic distance).

The embedding Φ is defined to be a critical point of E under boundary constraints that guarantee the image of Φ is indeed the designated hyperbolic orbifold $\mathcal{O}$. These boundary constraints are derived from the basic tile described in Section 3, as follows: First, we fix each Φ_j which corresponds to a landmark p_j to its corresponding cone position c_j in the basic tile, via a linear-equality constraint:

$$\Phi_j = c_j, \quad \forall j \in \mathcal{I}_c,$$

where $\mathcal{I}_c$ denotes the set of indices of all boundary vertices in $\mathbf{V}$ that are associated with landmarks. Second, for all other boundary vertices and their duplicates, $v_i, v_{i'}$, we require that for each pair, the two vertices are related by the associated isometry m_i , that is,

$$\Phi_{i'} = m_i(\Phi_i), \quad \forall (i, i') \in \mathcal{I}_b,$$

where the index set $\mathcal{I}_b = \{(i, i')\}$ contains the indices of all pairs of non-landmark boundary vertices in $\mathbf{V}$ associated to one another, and m_i is the isometry m associated to the arc on which Φ_i lies - refer to the inset and to Figure 4.

Finally, Φ is computed by finding a critical point of the optimization problem:

$$\min_{\Phi} E(\Phi) \quad (3a)$$

$$\text{s.t.} \quad \Phi_j = c_j, \quad j \in \mathcal{I}_c \quad (3b)$$


$$\Phi_{i'} = m_i(\Phi_i), \quad (i, i') \in \mathcal{I}_b \quad (3c)$$

We provide the full implementation details in Section 5. Figure 1 shows an example of a hyperbolic orbifold Tutte embedding of a sphere-type mesh of an owl into a sphere-type hyperbolic orbifold with 7 cones (the relevant basic tile is shown in Figure 4, left) computed via the optimization of (3).

By computing the Tutte embedding Φ of the cut mesh $\mathbf{M}^\circ$ we define a bijective, boundaryless, discrete-harmonic embedding of the mesh $\mathbf{M}$ into the orbifold $\mathcal{O}$, which is, by a slight abuse of notation, also denoted as Φ , that is $\Phi : \mathbf{M} \rightarrow \mathcal{O}$. This bijection is defined as follows: Every vertex u in $\mathbf{M}$ coincides with at least one vertex v_i in $\mathbf{M}^\circ$ which in turn is mapped to $\Phi_i \in D$. Now, define the mapping

$$\Phi(u) = [\Phi_i], \quad (4)$$

where $[\Phi_i]$ denotes the *orbit* of Φ_i , that is the set of all points $\{m(\Phi_i)\}$ reachable in D by transforming Φ_i using all the isometries from the orbifold symmetry group $\{m\}$. Note that each orbit $[z]$ corresponds to exactly one point of the orbifold according to the definition given in Section 3, and that $[\Phi(v_i)] = [\Phi(v_{i'})]$ in case $u = v_i$ is a vertex on the cut; hence $\Phi(u)$ is well-defined. In the next subsection we prove that the resulting map Φ is a bijection.



Theoretical properties Our goal in this section is to prove that a critical point Φ of the optimization problem (3) is guaranteed to yield a bijective embedding of $\mathbf{M}$ into $\mathcal{O}$. In the next subsection we will demonstrate how we practically find a critical point to this smooth optimization problem using standard optimization techniques.

First, let us establish a direct connection between the critical points of (3) and *Euclidean* convex combination maps, as defined by Floater [2003a]. This will be done using a result by Karcher [1977] that provides an expression for the gradients of the energy E :

$$\text{grad}_{\Phi_i} E(\Phi) = - \sum_{j \in \mathcal{N}_i} w_{ij} \exp_{\Phi_i}^{-1}(\Phi_j), \quad (5)$$

where $\text{grad}_{\Phi_i} E$ is the intrinsic gradient of E w.r.t. Φ_i , $\mathcal{N}_i$ is the index set of neighboring vertices to v_i , and $\exp_{\Phi_i}^{-1}(\Phi_j)$ is the inverse exponential map at Φ_i applied to Φ_j . Using (5) we can immediately prove a hyperbolic analog to the classic Tutte-embedding theorem,

Theorem 1. *Let $\mathbf{M}$ be a 3-connected disk-type surface. A critical point Φ of E that maps the boundary vertices of $\mathbf{M}$ to a convex polygon in the hyperbolic plane defines a bijective embedding.*

Before discussing the proof, we note that a convex polygon in the hyperbolic plane is defined similarly to the Euclidean case, namely, the unique geodesic between every pair of points in the polygon is contained within the polygon.

Theorem 1 is a direct consequence of the convex-combination-map theorem [Floater 2003a]: consider a critical Φ and map it into the Klein model K , where geodesics are straight lines (see Appendix A, and [Cannon et al. 1997] for the relation between the Poincaré model and the Klein model of the hyperbolic plane). Since Φ is critical, all intrinsic gradients are zero, $\text{grad}_{\Phi_i} E(\Phi) = 0$. Equation (5) then implies that every Φ_i is in *some* Euclidean strict convex combination of its neighbors. Lastly, hyperbolic convex polygons in Klein are also Euclidean-convex so the boundary is convex in the Euclidean sense.

We now move on to the orbifold case. We prove:

Theorem 2. *Let $\mathbf{M}$ be a 3-connected disk-type or sphere-type mesh with k landmarks. A critical solution to (3) defines a bijective embedding of $\mathbf{M}$ into the relevant orbifold $\mathcal{O}$.*

The idea behind the proof is to use the critical Φ to tile the entire hyperbolic plane using the Möbius and/or anti-Möbius transformations of the corresponding orbifold symmetry group. The key observation is that this provides an embedding $\hat{\Phi}$ of an infinite, simply-connected mesh $\hat{\mathbf{M}}$ such that all vertices are in the convex-hull of their one-ring neighbors. For interior and boundary (non-cone) vertices, this is a direct consequence of Eq. (5). For cone vertices, this is due to each cone being a kaleidoscopic or gyration point of the symmetry group [Conway et al. 2008]. We then use arguments in the spirit of Lovász [2004] and Floater [2003a] to prove that this implies the embedding is bijective. The full proof is detailed in Appendix B.

5 Implementation details

To compute a critical point of Problem (3) we use the Limited-memory BFGS (L-BFGS) algorithm, which is a first-order optimization method enabling linear constraints; we use the implementation provided in Matlab. The hyperbolic distance $d(z, w)$ in the Poincaré disk has an analytic expression (see Appendix A, Eq. 6) and is used to derive analytic expressions for the Dirichlet energy E and its (Euclidean) gradients $\nabla_{\Phi} E$. The linear part of the boundary conditions, that is (3b), is incorporated as-is into the optimization,


as linear-equality constraints. To deal with the non-linear boundary conditions (3c) we make the following observations for the sphere (or, generally, any boundaryless surface) and disk cases (or any surface with a boundary).

The spherical case Consider a boundary pair of vertices $v_i, v_{i'} \in \mathbf{V}$, $(i, i') \in \mathcal{I}_b$. We plug (3c) in E , cancelling $\Phi_{i'}$ from the optimization; the computation of the energy itself remains straightforward, and for the gradients we note that using the invariance of the hyperbolic distance to isometries we have:

$$\nabla_{\Phi_i} E = \sum_{j \in \mathcal{N}_i} w_{ij} \nabla_{\Phi_i} d(\Phi_i, \Phi_j) + \sum_{j \in \mathcal{N}_{i'}} w_{i'j} \nabla_{\Phi_i} d(\Phi_i, m_i^{-1}(\Phi_j))$$

which allows computing the gradient w.r.t. Φ_i . Practically, we noticed that preconditioning the gradients by the inverse metric, $(1 - |\Phi_i|^2)^2 \nabla_{\Phi_i} E$, increases the convergence speed.

The disk case In this case $v_i = v_{i'}$ and the isometry m_i is a hyperbolic reflection w.r.t. a geodesic line ℓ on which the relevant arc of the basic tile lies, refer also to Figure 4, right. Since m_i maps each point on ℓ to itself, the constraint (3c) simply entails that Φ_i should be constrained to ℓ . In this case it is beneficial to work in the Klein model K where geodesics are straight lines, as can be seen in the inset, which shows the same image from Figure 4, right, now in the Klein model. In K , the constraint (3c) becomes a linear equation, constraint (3b) stays linear, and the hyperbolic distance has an analytic expression which is similar to the Poincaré model (see Appendix A, Eq. (7)); as in the previous case, we derive analytic expressions for the energy value E and gradient $\nabla_{\Phi} E$. Note that since the Klein model, as opposed to the Poincaré model, does not possess a scaled Euclidean metric, we cannot employ the same preconditioner as in the spherical case.



Numerical stability and postprocessing A common pitfall of parameterization algorithms is numerical instability caused by large scale differences in the embedding. These occur, for example, when embedding a long perturbing part with no cone singularities. In our optimization scheme, this ill-conditioning manifests as differences in the magnitude of the per-vertex gradients, $\nabla_{\Phi_i} E$, between different vertices.

This is a relatively rare problem: out of the 190 embeddings computed for this paper only 5 exhibited this numerical instability; in these cases the standard L-BFGS terminates while some vertices, in highly shrunk areas, are still not in the hyperbolic convex hull of their neighbors. As this is a simple conditioning problem caused by considering highly different scales simultaneously, a simple solution is the following: when the output of L-BFGS is not injective in shrunk areas, perform gradient-descent independently per each 1-ring until convergence. This method has been employed in previous works [Hormann and Greiner 2000; Schreiner et al. 2004; Tsui et al. 2013; Shi et al. 2013]: for each vertex not in the hyperbolic convex-hull of its neighbors - use a Möbius transformation to translate the vertex to the origin along with its 1-ring (see Appendix C, Equation (8)), reposition the vertex to the Euclidean weighted centroid of its neighbors, and apply the inverse Möbius transformation to move it back into place. This process is repeated until all vertices are in the convex-hull of their 1-ring neighbors, which usually happens in 10-20 iterations.

Initialization To initialize the optimization we compute a standard *Euclidean* Tutte embedding and consider it as an embedding in the Klein model K . This is a feasible embedding into the relevant basic tile. The construction of the basic tile and boundary constraints for Equation (3) is rather technical and discussed in Appendix C. The embedding process is summarized in Algorithm 1.

Algorithm 1: Hyperbolic Orbifold Tutte embedding

input : Mesh $\mathbf{M}$, landmarks $p_1, \dots, p_k$;
Topologically-equivalent orbifold $\mathcal{O}$ with k cones.
output: Hyperbolic Tutte embedding, $\Phi : \mathbf{M} \rightarrow \mathcal{O}$.

- 1 Compute boundary conditions of the basic tile of $\mathcal{O}$ (Appendix C).
 - 2 Cut the mesh $\mathbf{M}$ to a disk mesh, $\mathbf{M}^\circ$ (Section 4).
 - 3 Compute Φ by optimizing Problem (3).
-

6 Evaluation

In this section we evaluate our algorithm, by experimenting with different choices of the weights w_{ij} for the Dirichlet energy (2), present tests regarding the effect of the cones on the embedding, and compare our method to another state-of-the-art method for hyperbolic embeddings.




Figure 5: Three embeddings into the same orbifold using different edge-weights, from left to right: combinatorial, mean-value, and cotan weights. As long as the weights are positive the embedding is guaranteed to be bijective.

Different weights As proven in Theorem 2, as long as the weights are positive the resulting embedding is guaranteed to be a bijection. In Figure 5 we compare embeddings into the same sphere-type orbifold with the same 9 prescribed cones for different choices of positive weights (from left to right): combinatorial weights; mean-value weights [Floater 2003b]; and cotan weights [Pinkall and Polthier 1993]. Evidently, the cotan weights yield an embedding which best preserves geometric details, however they are only positive for intrinsic Delaunay meshes. The mean-value weights are guaranteed to be positive, however they usually introduce more distortion into the embedding (*e.g.*, the ear of the bust in the blowups). Lastly, the combinatorial weights are derived only according to the connectivity of the mesh and hence the embedding can introduce an arbitrary amount of distortion. In the rest of the examples in this paper we used cotan weights and clamped all values smaller than $\varepsilon = 10^{-6}$ to ε to ensure positive weights; practically, this clamping had negligible effect on the embedding. An alternative solution is to subdivide the mesh to be intrinsic Delaunay [Fisher et al. 2007; Liu et al. 2015].

Number of cones Figure 6 depicts a stress-test of our algorithm’s robustness to varying the number of prescribed cones. We computed several embeddings of the same mesh with a varying number of cones, ranging from 5 and up to 50; in all cases, our algorithm produced a bijective embedding into the corresponding orbifold.




Figure 6: Embeddings of the same surface with an increasing number of cones, from 5 to 50.

Cone positioning The choice of which vertices of the surface are mapped to each cone of the target orbifold affects the resulting embedding. Figure 7 shows different embeddings of the same surface into the same 6-coned spherical orbifold. In the top row we show embeddings with the same 6 points on the bust chosen to be cones, but with different orderings; this already has a significant effect on the embedding, *e.g.*, in the top-right embedding, where an asymmetric ordering of the cones yields an asymmetric embedding. In the bottom row we show different, more arbitrary cone choices.




Figure 7: Embeddings into a 6-coned hyperbolic orbifold, each embedding resulting from a different selection of cones.

Disk orbifolds and convex polygons The target domain of a hyperbolic disk-orbifold embedding is a hyperbolic convex polygon; hence any embedding to a disk-orbifold can also be considered as an embedding to a hyperbolic convex polygon. In Figure 8 we compare the hyperbolic extension of the “classic” Tutte embedding of [Tutte 1963], where we minimize the Dirichlet energy while fixing the boundary map constraints, to the Hyperbolic Orbifold Tutte embedding into a disk orbifold which allows the boundary vertices (except for the cones) to slide on the edges of the target convex polygon. Theorems 1 and 2 (resp.) ensure that in both cases, minimizing the Dirichlet energy with these boundary conditions will result in a bijective embedding. As expected, allowing the boundary vertices to move yields an embedding with considerably lower Dirichlet energy.




Figure 9: Comparison of our method to CETM [Springborn et al. 2008]. Left: Our embedding (top) and CETM’s (bottom) of the same model with the same prescribed cones to a sphere orbifold (for CETM we show the generated double-cover). Right: 5 examples of cases where CETM failed to produce a valid bijection into the orbifold, while our method produces a bijection (our embedding shown next to each model). All meshes are intrinsic-Delaunay.




Figure 8: Two embeddings of the same disk mesh: once to a disk orbifold (left) and once to a convex polygon with a fixed boundary (right). Note that the disk orbifold embedding has a considerably lower Dirichlet energy (written at the bottom of each embedding).

Comparison to Conformal Equivalence of Triangle Meshes


CETM [Springborn et al. 2008] is a powerful method for computing discrete conformal embeddings of surfaces into Euclidean and hyperbolic domains with prescribed cone singularities. CETM can be used to compute embeddings into hyperbolic orbifolds by assigning cones and angles that enforce a hyperbolic orbifold structure. We compared our method to CETM using the author’s implementation². We show a comparison of our embedding to CETM’s in Figure 9, left. The two methods differ in several aspects: CETM’s embedding approximates a smooth conformal mapping while our embedding approximates harmonic mapping; CETM is not discrete-harmonic in each coordinate. Furthermore, while our algorithm embeds the surface into a predesignated orbifold, CETM embeds the surface into a conformally-equivalent orbifold - this means that two different surfaces will generally be mapped to two different orbifolds. Lastly, while CETM is a powerful and versatile algorithm

that can be applied in many different settings (e.g., both Euclidean and hyperbolic), it is not guaranteed to define a valid metric and hence may not yield a valid embedding. In comparison, the Hyperbolic Orbifold Tutte embedding is tailor-made for embedding into hyperbolic orbifolds and is guaranteed to be a bijection. In Figure 9, right, we show 5 intrinsic-Delaunay meshes for which CETM failed to find a valid embedding. Next to each mesh we show our algorithm’s bijective embedding.

7 Application to surface maps

Hyperbolic embeddings are useful for other applications, e.g., computation of surface maps [Tsui et al. 2013]. Hyperbolic Orbifold Tutte embeddings can be used to compute consistent maps between a collection of meshes.

Given two or more surfaces $\mathbf{M}_1, \mathbf{M}_2, \dots, \mathbf{M}_n$, $n \geq 2$, with k labeled points on each of the models, $p_{m,i} \in \mathbf{M}_m$, $i = 1, \dots, k$, we cut all surfaces to disks, $\mathbf{M}_m^\circ$, and compute the convex-combination maps $\Phi^m : \mathbf{M}_m^\circ \rightarrow D$, as detailed in Sections 4 and 5. In order to compute the surface map $f_{a,b} : \mathbf{M}_a \rightarrow \mathbf{M}_b$ between two arbitrary surfaces in the collection, $1 \leq a, b \leq n$, we use Equation (4): given a vertex v_i in $\mathbf{M}_a^\circ$ we search an element in the orbit of its image, $[\Phi_i^a]$, that is contained in a hyperbolic triangle of $\Phi^b(\mathbf{M}_b^\circ)$. In the inset we show two Hyperbolic orbifold Tutte embeddings overlaid on one another. The Black Star, Φ_i^a , belonging to the red embedding Φ^a , is not in the image of any triangle of the blue embedding Φ^b , however the white star, $m(\Phi_i^a)$, is the unique point in the orbit $[\Phi_i^a]$ which lies in a triangle of the blue embedding Φ^b . Practically, we search the orbit $[\Phi_i^a]$ by repeatedly applying a transformation m from the orbifold symmetry group to the point Φ_i^a and checking whether it lands inside a triangle of $\Phi^b(\mathbf{M}_b^\circ)$.



²Varylab, available at www.varylab.com.



Figure 10: (a) A mapping between a pair of meshes from Figure 2, shown from seven different viewpoints; the map is guaranteed to be a homeomorphism between the two surfaces. Note that we re-texture the source mesh from each viewpoint. In (b) we show another mapping from several viewpoints.

As all surface mappings are defined via $f_{a,b} = (\Phi^b)^{-1} \circ \Phi^a$ (remember that by Φ also denotes the surface-to-orbifold map, see Section 4), they are cycle-consistent in the following sense [Nguyen et al. 2011]:

$$f_{b,c} \circ f_{a,b} = (\Phi^c)^{-1} \circ \Phi^b \circ (\Phi^b)^{-1} \circ \Phi^a = f_{a,c}.$$

That is, composition of the mappings $\mathbf{M}_a \rightarrow \mathbf{M}_b$ and $\mathbf{M}_b \rightarrow \mathbf{M}_c$ coincides with the mapping $\mathbf{M}_a \rightarrow \mathbf{M}_c$. This property is important in analyzing a collection of shapes as it is a necessary condition for well-defining correspondences, textures and labels across a collection of surfaces.

The last remaining issue is choosing the cutting order of the landmarks $p_{m,i}$ of each mesh $\mathbf{M}_m$ (in case the meshes are spherical and need to be cut-open to disks); this is equivalent to choosing which landmark $p_{m,i}$ is mapped to each cone of $\mathcal{O}$. Clearly, this entails that the *order* of the cuts should be consistent in all meshes $\mathbf{M}_m$ in order to achieve correct interpolation of the landmarks. In addition to the order of the cut, the cut itself sets a homotopy class for the mappings between the surfaces and the orbifold $\mathcal{O}$. We found that a good heuristic for setting both the order of the points and the cuts themselves is to find an approximate solution to the Traveling Salesman Problem (TSP) defined for one of the surfaces over a graph with the landmarks as vertices and the landmarks' geodesic distances assigned as edge-weights between all pairs of landmarks. We use a simple TSP approximation algorithm [Miller et al. 1960] which yields nicely-behaved cuts in all cases we experimented on.

Visualization of surface maps To visualize the surface-to-surface maps we texture one mesh (the source mesh), and use the map to transfer the texture to the other meshes. To achieve a high quality, low-distortion texture we use the same method used in [Aigerman et al. 2014; Aigerman et al. 2015] of only texturing the parts visible from the camera's viewpoint. Figure 10 shows two of the computed surface maps from different viewing directions, visualized by rotating the models and generating a different texture on the source mesh per viewing direction.

Mapping spherical surfaces We used our algorithm to compute collective mappings of the different shape-classes from the dataset of SHREC07 [Giorgi et al. 2007], using the landmarks created in [Kim et al. 2011]. We computed collective mappings on all models, while excluding cases with topological problems. In Figure 11 we show the collective maps within the humans class and the hands class. The mappings naturally accommodate for the different proportions of the different humans, e.g., on the boy's head (second row, third from right). Furthermore, note that while the boy is not isometric to most of the other models, it is quite isometric to the girl (top row, third from left). Although they are both somewhat outliers to the rest of the set, our method works well in this case and indeed the computed map between the two has low isometric distortion, e.g., on their heads. Likewise, in the hands collection, maps between similar models like the synthetic hands (most of the top row) exhibit high accuracy. Less isometric pairs, for example the elongated, bottom rightmost hand, and the bottom leftmost hand are also mapped naturally. In Figure 12 we show mappings of the busts and four-legged animals. Although the busts collection only has a few landmarks on the faces of the models, the resulting maps still present an accurate mapping of facial features. While the bodies of the busts have no landmarks marked on them, the produced maps elegantly handle the change in size and shape of each model. The algorithm produces natural maps for the four-legged animals collection containing models which vary significantly from one-another by shape and posture. The giraffe is an example of a less natural mapping produced by our method - the elongated neck presents distortion which is potentially avoidable when considering a mapping of *one* other animal to the giraffe, and not a collection.

Mapping disk domains Embeddings to disk-type orbifolds can be useful to compute mappings between disk-type domains. The advantage of this method is that the boundary is free to move (aside from any selected landmarks) along the orbifold's boundary during optimization, and as a result the boundary map between the two meshes is also optimized. In Figure 13 we show mappings of two collections of planar domains of silhouettes. The ability of the boundary map to change is evident near the rightmost woman's elongated hair. Likewise, for the dog models who have very different shapes (e.g., different number of visible legs, open and closed mouths) our mapping still produces plausible results.




Figure 11: Collective mappings computed on two collections from the SHREC07 [Giorgi et al. 2007] dataset.

7.1 Relation to other surface-mapping methods

Globally optimal cortical surface matching [Tsui et al. 2013]

Tsui et al. [2013] were the first to suggest using spherical orbifolds as common domains in order to map cortical surfaces. Their method first uses CETM [Springborn et al. 2008] to embed the meshes into conformally-equivalent orbifolds and then in a second step approximates a harmonic mapping between the orbifolds. As the comparison discussed in Section 6 shows, CETM can fail to provide a bijective mapping in many instances; in all such cases, there is no designated orbifold structure and the method of [Tsui et al. 2013] cannot provide a bijective surface-map. The inset shows a map computed via our algorithm between two octopus models; since CETM failed to embed the octopus, the method of [Tsui et al. 2013] cannot compute a mapping of the two models. A key observations of the current paper is that there exist closed-form boundary conditions, extracted from the hyper-



bolic symmetry groups, that together with Tutte's embedding can be proven to yield bijective, discrete harmonic orbifold embeddings. Other differences between the methods are: (i) the current paper does not require the conformal parameterization step performed in Tsui et al. as it directly maps the two surfaces into the *same* orbifold; (ii) since CETM is not discrete harmonic, the maps computed in Tsui et al. are not discrete harmonic; (iii) other orbifold types, e.g., disks, which require different treatment, are considered in this work; (iv) the current paper formulates the problem as an unconstrained, smooth optimization problem.

Locally injective parameterization [Weber and Zorin 2014]

Weber and Zorin suggest a method to map disk domains by computing Euclidean Tutte embeddings of two surfaces into a common domain. We compare our method to theirs in Figure 14. Our method computes the surface map via two Tutte embeddings into a disk orbifold, and hence the boundary map can change during optimization so as to lower the harmonic energy further. Weber et al. do not optimize the boundary map, which can produce higher shear and distortion as shown in the blowups (See also Figure 8).



Figure 12: Two more collective mappings of collections from the SHREC07 data-set.

Seamless surface mappings [Aigerman et al. 2015] Aigerman *et al.* also compute a surface map by embedding the two meshes into a cone manifold. They can also map small collections, however since all models need to be incorporated into a single, large second-order cone program their method scales roughly cubically in the number of models and in practice cannot map more than a handful of models (they report half an hour for 3 medium-size meshes). Since our method uses a predefined, fixed canonical domain, the embedding of each mesh is decoupled and our computation scales linearly with the number of models. In Figure 15 we compare the quality of the map we computed, versus the one computed using [Aigerman et al. 2015]. We produce overall comparable results; note however, that the symmetric structure of the orbifold alleviates the harsh pinching near the landmarks evident in the result of Aigerman *et al.*

7.2 Technical details

The algorithm was implemented in Matlab. Typical timings of our algorithm are detailed in Table 1. The computations were performed on a single thread on a 3.50GHz Intel i7 CPU. The surface-

map computation can be significantly sped by using, *e.g.*, KD-trees for the point-in-triangle lookup.

Note that for the disk models the optimization of Equation (3) is slower, since in the disk-orbifold case we perform the optimization in the Klein model which is less well-behaved (*i.e.*, more scaled and anisotropic) than the Poincaré model and does not possess similar preconditioning (see Section 5).

8 Conclusions

Tutte’s embedding is extended to the hyperbolic case and is proven to yield bijective maps to orbifolds and convex polygons. The family of hyperbolic orbifolds is extremely rich, and it seems worthwhile to explore more uses of different orbifolds in the future. For example, our method can be readily applied to higher genus surfaces - the only difference is in the cutting mechanism; using Tutte’s embedding on higher genus models without any cones could yield an alternative uniformization technique.

The ability to efficiently compute consistent homeomorphisms in large collections of surfaces opens up numerous future research



Figure 13: Collective mappings computed for two collections of planar disk-type domains.

directions such as shape-space exploration and synthesis of new shapes; shape-space statistics and distance computation; computation of a mean-shape of a collection; and globally-consistent meshing of a set of surfaces.

Our method has a few shortcomings. Using a fixed common domain enables fast and robust computation of collective mappings, however it may hinder quality, *e.g.*, in the Giraffe model in Figure 12. Nonetheless, the vast majority of the maps are of high quality, and all maps are guaranteed to be seamless homeomorphisms.

The second shortcoming of our method in case cuts need to be introduced (such as the spherical case) is that it relies on the way the cut is chosen (for an elaborate discussion see [Aigerman et al. 2015], Section 8, paragraph “Homotopy classes”). While in practice our cutting method works well, cutting along the shortest path is only a heuristic. Defining what is the “correct” homotopy class of the mappings and generating cuts which are in this homotopy class will make this algorithm - as well as other methods which rely on these type of cuts - more robust, and we plan to investigate this problem in the future.

The last limitation of our method is that it is not fully automatic; the choice of cone-vertices (their number and their location on the mesh) is currently supplied by the user. Optimal selection of cones so as to reduce distortion will result in higher-quality embeddings, as can be seen in Figure 7; automating this task is also one of our future goals.

9 Acknowledgements

This work was funded by the European Research Council (ERC Starting Grant SurfComp), the Israel Science Foundation (grant No. 1284/12) and the I-CORE program of the Israel PBC and ISF (Grant No. 4/11). The authors would like to thank Boris Springborn and Stefan Sechelmann for their implementation of CETM in Varylab, Raz Slutsky for proofreading, and the anonymous reviewers for their comments and suggestions. All meshes are from the datasets mentioned within the text, except for the owl, elder and yeti which are from Pinshape.

Model	#V	#T	#p	Init	Tutte	Map
FAUST	7K	14K	23	4	41	42
Humans	10K	20K	36	6	45	60
Hands	8K	14K	17	9	46	17
Busts	27K	53K	11	9	189	57
Animals	8K	14K	22	4	16	28
2D	11K	19K	9	1	509	1
Owl	25K	50K	7	7	301	-
Igea	3K	6K	7	1	12	-
Yeti	25K	50K	51	24	314	-
Elder (MVC)	20K	40K	9	6	510	-
Elder (Cot)					300	-
Elder (Comb)					80	-

Table 1: Timings (in seconds) of our algorithm on a member from each collection (top) and other embeddings (bottom). In each row we detail the number of vertices, faces, landmarks, initialization time (mainly cutting the mesh and initial bookkeeping), embedding time, and computing the map to another model in the class.

References

- AIGERMAN, N., AND LIPMAN, Y. 2015. Orbifold tutte embeddings. *ACM Trans. Graph.* 34, 6 (Oct.), 190:1–190:12.
- AIGERMAN, N., PORANNE, R., AND LIPMAN, Y. 2014. Lifted bijections for low distortion surface mappings. *ACM Trans. Graph.* 33, 4 (July), 69:1–69:12.
- AIGERMAN, N., PORANNE, R., AND LIPMAN, Y. 2015. Seamless surface mappings. *ACM Transactions on Graphics (TOG)* 34, 4, 72.
- BOGO, F., ROMERO, J., LOPER, M., AND BLACK, M. J. 2014. FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Piscataway, NJ, USA.
- BOMMES, D., ZIMMER, H., AND KOBELT, L. 2009. Mixed-integer quadrangulation. *ACM Trans. Graph.* 28, 3 (July), 77:1–77:10.




Figure 14: Our result (right), and [Weber and Zorin 2014] (left). Source model on the top. Allowing the boundary map to move in our method reduces the distortion and yields a smooth map.

- BRADLEY, D., POPA, T., SHEFFER, A., HEIDRICH, W., AND BOUBEKEUR, T. 2008. Markerless garment capture. *ACM Trans. Graph.* 27, 3 (Aug.), 99:1–99:9.
- CAMPEN, M., BOMMES, D., AND KOBELT, L. 2015. Quantized global parametrization. *ACM Transactions on Graphics (TOG)* 34, 6, 192.
- CANNON, J. W., FLOYD, W. J., KENYON, R., PARRY, W. R., ET AL. 1997. Hyperbolic geometry. *Flavors of geometry* 31, 59–115.
- CONWAY, J. H., BURGIEL, H., AND GOODMAN-STAUBS, C. 2008. *The symmetries of things*. A.K. Peters, Wellesley (Mass.).
- DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. Intrinsic parameterizations of surface meshes. In *Computer Graphics Forum*, vol. 21, Wiley Online Library, 209–218.
- FISHER, M., SPRINGBORN, B., SCHRÖDER, P., AND BOBENKO, A. I. 2007. An algorithm for the construction of intrinsic delaunay triangulations with applications to digital geometry processing. *Computing* 81, 2-3, 199–213.
- FLOATER, M. 2003. One-to-one piecewise linear mappings over triangulations. *Mathematics of Computation* 72, 242, 685–696.
- FLOATER, M. S. 2003. Mean value coordinates. *Computer Aided Geometric Design* 20, 1, 19–27.
- GIORGI, D., BIASOTTI, S., AND PARABOSCHI, L. 2007. SHREC: SHape REtrieval Contest: Watertight models track. <http://watertight.ge.imati.cnr.it/>.




Figure 15: Comparison of our algorithm (right) to [Aigerman et al. 2015] (middle). Source model shown on the left. At the bottom row we show blowups exhibiting the pinching at the landmarks, which is alleviated in our result.

- GORTLER, S. J., GOTSCHMAN, C., AND THURSTON, D. 2006. Discrete one-forms on meshes and applications to 3d mesh parameterization. *Computer Aided Geometric Design* 23, 2, 83–112.
- HORMANN, K., AND GREINER, G. 2000. MIPS: An efficient global parametrization method. In *Curve and Surface Design: Saint-Malo 1999*, P.-J. Laurent, P. Sablonnière, and L. L. Schumaker, Eds., Innovations in Applied Mathematics. Vanderbilt University Press, Nashville, TN, 153–162.
- HORMANN, K., LÉVY, B., AND SHEFFER, A. 2007. Mesh parameterization: Theory and practice video files associated with this course are available from the citation page. In *ACM SIGGRAPH 2007 Courses*, ACM, New York, NY, USA, SIGGRAPH '07.
- JIN, M., KIM, J., LUO, F., AND GU, X. 2008. Discrete surface Ricci flow. *Visualization and Computer Graphics, IEEE Transactions on* 14, 5, 1030–1043.
- KARCHER, H. 1977. Riemannian center of mass and mollifier smoothing. *Communications on pure and applied mathematics* 30, 5, 509–541.
- KHAREVYCH, L., SPRINGBORN, B., AND SCHRÖDER, P. 2006. Discrete conformal mappings via circle patterns. *ACM Transactions on Graphics (TOG)* 25, 2, 412–438.
- KIM, V. G., LIPMAN, Y., AND FUNKHOUSER, T. 2011. Blended intrinsic maps. *ACM Trans. Graph.* 30, 4 (July), 79:1–79:12.
- KRAEVOY, V., AND SHEFFER, A. 2004. Cross-parameterization and compatible remeshing of 3d models. *ACM Trans. Graph.* 23, 3 (Aug.), 861–869.
- LEE, A. W. F., DOBKIN, D., SWELDENS, W., AND SCHRÖDER, P. 1999. Multiresolution mesh morphing. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, SIGGRAPH '99, 343–350.
- LIPMAN, Y. 2012. Bounded distortion mapping spaces for triangular meshes. *ACM Trans. Graph.* 31, 4 (July), 108:1–108:13.
- LIU, Y.-J., XU, C.-X., FAN, D., AND HE, Y. 2015. Efficient construction and simplification of delaunay meshes. *ACM Transactions on Graphics (TOG)* 34, 6, 174.

LOVÁSZ, L. 2004. Discrete analytic functions: an exposition. *Surveys in differential geometry* 9, 241–273.

MILLER, C. E., TUCKER, A. W., AND ZEMLIN, R. A. 1960. Integer programming formulation of traveling salesman problems. *Journal of the ACM (JACM)* 7, 4, 326–329.

MYLES, A., AND ZORIN, D. 2013. Controlled-distortion constrained global parametrization. *ACM Trans. Graph.* 32, 4 (July), 105:1–105:14.

NGUYEN, A., BEN-CHEN, M., WELNICKA, K., YE, Y., AND GUIBAS, L. 2011. An optimization approach to improving collections of shape maps. *Computer Graphics Forum* 30, 5, 1481–1491.

PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 15–36.

SCHREINER, J., ASIRVATHAM, A., PRAUN, E., AND HOPPE, H. 2004. Inter-surface mapping. *ACM Trans. Graph.* 23, 3 (Aug.), 870–877.

SCHÜLLER, C., KAVAN, L., PANOZZO, D., AND SORKINE-HORNUNG, O. 2013. Locally injective mappings. In *Computer Graphics Forum*, vol. 32, Wiley Online Library, 125–135.

SHEFFER, A., LÉVY, B., MOGILNITSKY, M., AND BOGOMYAKOV, A. 2005. Abf++: fast and robust angle based flattening. *ACM Transactions on Graphics (TOG)* 24, 2, 311–330.

SHEFFER, A., PRAUN, E., AND ROSE, K. 2006. Mesh parameterization methods and their applications. *Found. Trends. Comput. Graph. Vis.* 2, 2 (Jan.), 105–171.

SHI, R., ZENG, W., SU, Z., DAMASIO, H., LU, Z., WANG, Y., YAU, S.-T., AND GU, X. 2013. Hyperbolic harmonic mapping for constrained brain surface registration. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, 2531–2538.

SMITH, J., AND SCHAEFER, S. 2015. Bijective parameterization with free boundaries. *ACM Trans. Graph.* 34, 4 (July), 70:1–70:9.

SPRINGBORN, B., SCHRÖDER, P., AND PINKALL, U. 2008. Conformal equivalence of triangle meshes. *ACM Transactions on Graphics (TOG)* 27, 3, 77.

STEPHENSON, K. 2005. *Introduction to circle packing: The theory of discrete analytic functions*. Cambridge University Press.

TSUI, A., FENTON, D., VUONG, P., HASS, J., KOEHL, P., AMENTA, N., COEURJOLLY, D., DECARLI, C., AND CARMICHAEL, O. 2013. Globally optimal cortical surface matching with exact landmark correspondence. In *Proceedings of the 23rd International Conference on Information Processing in Medical Imaging*, Springer-Verlag, Berlin, Heidelberg, IPMI’13, 487–498.

TUTTE, W. T. 1963. How to draw a graph. *Proc. London Math. Soc* 13, 3, 743–768.

WEBER, O., AND ZORIN, D. 2014. Locally injective parametrization with arbitrary fixed boundaries. *ACM Transactions on Graphics (TOG)* 33, 4, 75.

Appendix A Hyperbolic models

The geodesic distance between two points in the Poincaré disk D is given by the formula

$$d(z, w) = \operatorname{arccosh} \left(1 + 2 \frac{|z - w|^2}{(1 - |z|^2)(1 - |w|^2)} \right). \quad (6)$$

The orientation-preserving isometries are the set of all Möbius transformations mapping the unit disk to itself. These transformations can be written explicitly for any point $z \in D$ as

$$\left\{ z \mapsto e^{i\theta} \frac{z + c}{1 + \bar{c}z} \mid c \in \mathbb{C}, \theta \in [0, 2\pi) \right\}.$$

The rest of the isometries of $\mathbb{H}^2$ - the anti-Möbius transformations - are the set comprised of each of the above transformations composed with the complex conjugation.

For a point $z \in \mathbb{C}$ in the Poincaré model, the corresponding point in the Klein model is given by $\frac{2z}{1+|z|^2}$. The distance between two points in the Klein model is given by

$$d(u, v) = \operatorname{arccosh} \left(\frac{1 - \operatorname{Re}(u\bar{v})}{\sqrt{1 - |u|^2} \sqrt{1 - |v|^2}} \right). \quad (7)$$

Appendix B Bijectivity of the Hyperbolic Orbifold Tutte embedding

In this appendix we present the sketch of the proof of Theorem 2. We consider $\mathbf{M}^\circ$ and its embedding Φ defined as a critical point of (3). We will use the Klein disk K as the target domain of Φ , that is $\Phi : \mathbf{V} \rightarrow K$. We use the Klein model in the proof in order to reuse lemmas from the Euclidean case. We will also consider the infinite mesh, $\hat{\mathbf{M}} = (\hat{\mathbf{V}}, \hat{\mathbf{E}}, \hat{\mathbf{T}})$, created by stitching copies of $\mathbf{M}^\circ$ according to the transformations in the orbifold group, and $\hat{\Phi}$ its mapping into the Klein disk K constructed from Φ .

The proof follows several steps: 1) showing that Φ does not degenerate at least one triangle of $\mathbf{M}^\circ$; 2) for every two triangles in $\hat{\mathbf{M}}$ sharing an edge, if $\hat{\Phi}$ does not degenerate one of the triangles, then it does not degenerate also the other triangle, and the images under $\hat{\Phi}$ of the two triangles lie on different sides of the edge; and 3) If all triangles mapped by $\hat{\Phi}$ are non-degenerate and have consistent orientation then $\hat{\Phi}$ is a bijective embedding onto K . This in particular implies that $\Phi : \mathbf{M} \rightarrow \mathcal{O}$ is bijective.

We start with (1). Assume that all triangles are degenerate, and consider an arbitrary triangle. Since it is degenerate it is contained in an infinite geodesic line. Its neighboring triangles are also assumed to be degenerate and have two points on that line, therefore they are also contained in the line. Continuing in this manner leads to all triangles contained in an infinite geodesic line which contradicts (3b). For (2) we will use a lemma (inspired by a part of the proof of Theorem 19 in [Lovász 2004] and [Floater 2003a]):

Lemma 1. *Consider an infinite geodesic line ℓ , and assume images of two vertices $\hat{\Phi}_i, \hat{\Phi}_j$ of $\hat{\mathbf{M}}$ are on one side of ℓ . Then, there exists a path of vertices connecting these two vertices that is also on the same side of ℓ .*

This Lemma is a bit technical and proved below. Let us use it to prove (2). Consider the setting described in (2): Let e_{ij} be an edge in $\hat{\mathbf{E}}$ and $t_{ijk}, t_{ijm} \in \hat{\mathbf{T}}$ two adjacent triangles. Without losing generality assume t_{ijk} is non-degenerate and that $\hat{\Phi}_k$ lies on one side of ℓ . Since we are now in the Klein disk, every vertex $\hat{\Phi}_i$ is in some Euclidean strict convex-hull of its one-ring neighbors (see explanation at the end of Section 4). Similar to [Floater

2003a] we define a linear function f that is zero on the infinite line (geodesic) containing $\widehat{\Phi}_i, \widehat{\Phi}_j$ and strictly positive on $\widehat{\Phi}_k$, that is $f(\widehat{\Phi}_k) > 0$. Due to the convex-hull property at $\widehat{\Phi}_i$ it has a neighbor such that $f(\widehat{\Phi}_{i'}) < 0$. Similarly, $\widehat{\Phi}_j$ has a neighbor $\widehat{\Phi}_{j'}$ such that $f(\widehat{\Phi}_{j'}) < 0$. Using lemma 1, we connect $\widehat{\Phi}_{i'}, \widehat{\Phi}_{j'}$ using a path P with vertices satisfying $f(\widehat{\Phi}_s) < 0$. Now consider the closed path P' defined by P concatenated with the path $\widehat{\Phi}_{i'} \rightarrow \widehat{\Phi}_i \rightarrow \widehat{\Phi}_j \rightarrow \widehat{\Phi}_{j'}$. We get a closed path with boundary values $f(\widehat{\Phi}_s) < 0$ except $f(\widehat{\Phi}_i) = 0 = f(\widehat{\Phi}_j)$. Considering the vertex $\widehat{\Phi}_m$ of the triangle ijm , then if $f(\widehat{\Phi}_m) \geq 0$, the discrete maximum principle (using the 3-connectedness) for convex combination maps (see [Floater 2003a]) implies f is non-negative on the entire path P' which contradicts that $f(\widehat{\Phi}_{i'}) < 0$. So $f(\widehat{\Phi}_m) < 0$ and (2) is proven.

To prove (3) we will use similar arguments to [Lipman 2012]. From (1) we have that there is one triangle in $\widehat{\Phi}$ that is not degenerate. We can now use (2) and deduce that its neighbors are also non-degenerate and also that they have the same orientation. We can continue in this manner and conclude that all triangles in $\widehat{\Phi}$ have the same orientation. To prove that $\widehat{\Phi}$ is injective and onto K we can count how many triangles contain a fixed arbitrary generic point $z \in K$. If we show there is exactly one, we are done. To check if z is in some triangle we can compute the winding number of z w.r.t. the boundary of the triangle; if it is inside we get ± 1 , depending on the orientation of the triangle, and if z is outside we will get 0. Since all triangles have the same orientation, the absolute value of the sum of winding numbers equals the number of preimages of z . Since Φ is contained in some compact subset of the Klein disk there is only a finite number of triangles that can possibly contain z . We can tile enough copies of Φ around z and ensure that z has a winding number 1 w.r.t. the boundary of this tiled piece and no other copy of Φ contains z . Since the winding number w.r.t. the boundary equals the sum of the winding numbers w.r.t. the individual triangles in the tiled piece - z is contained in exactly one triangle from $\widehat{\Phi}$.

This implies that Φ , considered as a map from $\mathbf{M}$ to $\mathcal{O}$ defined in (4), is bijective. Indeed, it is onto $\mathcal{O}$ since for every point $z \in K$ we can find an isometry m from the orbifold isometry group and a point $p \in \mathbf{M}^\circ$ such that $z = m(\Phi(p))$, since $\widehat{\Phi}$ is onto K . This implies that $\Phi(p) \in [z]$ where the orbit $[z]$ represents an arbitrary point in the orbifold $\mathcal{O}$. Φ is injective since otherwise there exist two points $p, q \in \mathbf{M}^\circ$, $p \neq q$ in $\mathbf{M}$, such that $[\Phi(p)] = [\Phi(q)]$ but that would imply that there exist two isometries m_1, m_2 from the orbifold group such that $m_1(\Phi(p)) = m_2(\Phi(q))$ which contradicts the injectivity of $\widehat{\Phi}$. $\square$

To complete the proof we need to prove Lemma 1. For that, we will first prove another lemma, Inspired again by the proof of Theorem 19 in [Lovász 2004],

Lemma 2. *Let ℓ be an infinite geodesic line in the hyperbolic plane, and $\widehat{\Phi}_i$ a vertex image not on the line. Then, one can find an infinite path $\widehat{\Phi}_{i1} \rightarrow \widehat{\Phi}_{i2} \rightarrow \dots$ emanating from $\widehat{\Phi}_i$ that stays on the same side of ℓ as $\widehat{\Phi}_i$ and the distance from ℓ along the path grows to infinity.*

We continue to work in the Klein disk K . We note that since geodesic lines are straight lines, hyperbolic convex sets in this model are Euclidean-convex and vice-versa. Without loss of generality we assume ℓ is the line $\{i \cdot t \mid -1 < t < 1\}$, and $\text{Re } \widehat{\Phi}_i > 0$. Denote $K_+ = K \cap \{z \mid \text{Re } z > 0\}$. Let us consider the submesh $\mathbf{W} \subset \widehat{\mathbf{M}}$ containing all vertices in $\widehat{\mathbf{M}}$ such that their images are reachable from $\widehat{\Phi}_i$ by paths that stay in K_+ . Let $U = K_+ \setminus \widehat{\Phi}(\mathbf{W})$

be the unreachable set from $\widehat{\Phi}_i$. Since the orbifold hyperbolic groups are discrete, their orbits do not have accumulation points in K . This means that U has a polygonal boundary. Furthermore, no exterior angles of this boundary polygon are concave since otherwise there exists a vertex not in the convex hull of its neighbors. This means that U is convex. Towards contradiction, assume there exists some constant $c > 0$ such that U contains all points in K with distance at-least c from ℓ . Consider the line $\ell' = (\ell + d) \cap K$, with $d > 0$. Near the ends of ℓ' the distance to ℓ is unbounded so there are points arbitrarily close to the ends of ℓ' that are in U , therefore due to convexity $\ell' \subset U$. Hence, $K_+ \subset U$ and in particular $\widehat{\Phi}_i \in U$, which is a contradiction. $\square$

Lastly, we finish the proof of Lemma 1. First, we use Lemma 2 to construct two paths P_i, P_j emanating from $\widehat{\Phi}_i, \widehat{\Phi}_j$ (resp.) with each path consisting of vertices from $\widehat{\Phi}$ so that when traversing the path the hyperbolic distance from ℓ goes to infinity. Since Φ has a finite number of vertices, it has a finite hyperbolic diameter $d > 0$. Let $\widehat{\Phi}_{i'} \in P_i$ and $\widehat{\Phi}_{j'} \in P_j$ be points of distance $4d$ to ℓ . Since $\mathbf{M}^\circ$ is connected we can connect $\widehat{\Phi}_{i'}$ and $\widehat{\Phi}_{j'}$ to the nearest cone in their respective tile. This path remains far from ℓ by at least a hyperbolic distance of $3d$. Now consider the hyperbolic grid created by the orbits of the cone points c_i of the basic tile and take all such point of hyperbolic distance greater than $2d$ to the line ℓ . Each two cones in this set can be connected with a path which is at least d far from ℓ , and in particular the two cones found above. Concatenating these five paths provides a path in K_+ . $\square$

Appendix C Construction of the basic tile

We now detail how we construct the boundary constraints used to compute the orbifold embedding Φ . We will use the fact that in the Poincaré model, Euclidean rotations around the origin are isometries. A Möbius transformation translating a point $p \in \mathbb{C}$ to the origin is defined via

$$m_p(z) = \frac{z - p}{1 - \bar{p}z} \quad (8)$$

If a point has hyperbolic distance l to the origin, then the geodesic of length l connecting it to the origin in the Poincaré model is a straight line, and its Euclidean length is $\frac{e^l - 1}{e^l + 1}$.

Disk-type basic tile We start with constructing the basic tile of the disk orbifold, that is a hyperbolic k -gon with $k \geq 5$ (e.g., Figure 4, right). Given angles $\theta_1, \theta_2, \theta_3, \dots, \theta_k$, $\sum \theta_i < \pi$, there exists a unique hyperbolic triangle (up to an isometry) with these three angles. Its edge lengths (l_1, l_2, l_3) are given by the formula

$$l_i = \text{arccosh} \left(\frac{\cos(\theta_i) + \cos(\theta_{i+1}) \cos(\theta_{i+2})}{\sin(\theta_{i+1}) \sin(\theta_{i+2})} \right),$$

where the indices are cyclically shifted. By dissecting the k -gon into k equilaterals through the origin, each of angles $\pi/4, \pi/4$ and $2\pi/k$ we get that the corners of the k -gon are of hyperbolic distance

$$l = \text{arccosh} \left(\frac{1 + \cos \frac{2\pi}{k}}{\sin \frac{2\pi}{k}} \right)$$

to the origin and therefore the corners of the k -gon are

$$p_j = e^{i \frac{2\pi j}{k}} \frac{e^l - 1}{e^l + 1}, \quad j = 1, \dots, k.$$

Sphere-type basic tile To construct the basic tile of the sphere orbifold, we use a Möbius transformation to translate the polygon so the geodesic midpoint between p_k and p_1 lies at the origin and rotate so that p_1, p_k are on the real axis. We now vertically reflect the polygon by applying a complex conjugation. By this we get a second copy of the k -gon, $q_1, q_2, \dots, q_k$ where $p_1 = q_1$ and $p_k = q_k$. These two polygons together define the basic tile of the sphere orbifold, with the arcs $p_j, p_{j+1}, j = 1, \dots, k-1$, associated to the arcs q_j, q_{j+1} , respectively.