A Pseudopolynomial Algorithm for Alexandrov's Theorem

Daniel Kane[1]*, Gregory N. Price[2]**, and Erik D. Demaine[3]***

[1] Department of Mathematics, Harvard University,
1 Oxford Street, Cambridge MA 02139, USA, dankane@math.harvard.edu

[2] MIT Computer Science and Artificial Intelligence Laboratory,
32 Vassar Street, Cambridge, MA 02139, USA, {price,edemaine}@mit.edu

Abstract. Alexandrov's Theorem states that every metric with the global topology and local geometry required of a convex polyhedron is in fact the intrinsic metric of some convex polyhedron. Recent work by Bobenko and Izmestiev describes a differential equation whose solution is the polyhedron corresponding to a given metric. We describe an algorithm based on this differential equation to compute the polyhedron to arbitrary precision given the metric, and prove a pseudopolynomial bound on its running time.

1 Introduction

Consider the intrinsic metric induced on the surface M of a convex body in $\mathbb{R}^3$. Clearly M under this metric is homeomorphic to a sphere, and locally convex in the sense that a circle of radius r has circumference at most $2\pi r$.

In 1949, Alexandrov and Pogorelov [1] proved that these two necessary conditions are actually sufficient: every metric space M that is homeomorphic to a 2-sphere and locally convex can be embedded as the surface of a convex body in $\mathbb{R}^3$. Because Alexandrov and Pogorelov's proof is not constructive, their work opened the question of how to produce the embedding given a concrete M.

To enable computation we require that M be a polyhedral metric space, locally isometric to $\mathbb{R}^2$ at all but n points (vertices). Now the theorem is that every polyhedral metric, a complex of triangles with the topology of a sphere and positive curvature at each vertex, can be embedded as an actual convex polyhedron in $\mathbb{R}^3$. This case of the Alexandrov-Pogorelov theorem was proven by Alexandrov in 1941 [1], also nonconstructively. Further, Cauchy showed in 1813 [3] that such an embedding must be unique. All the essential geometry of the general case is preserved in the polyhedral case, because every metric satisfying the general hypothesis can be polyhedrally approximated.

Algorithms for Alexandrov's Theorem are motivated by the problem of folding a polygon of paper into precisely the surface of a convex polyhedron. There

* Partially supported by an NDSEG Fellowship.
** Partially supported by an NSF Graduate Research Fellowship.
*** Partially supported by NSF CAREER award CCF-0347776.

are efficient algorithms to find one or all gluings of a given polygon's boundary to itself so that the resulting metric satisfies Alexandrov's conditions [4, 9]. But this work leaves open how to find the actual 3D polyhedra that can be folded from the polygon of paper.

In 1996, Sabitov [12, 11, 13, 5] showed how to enumerate all the isometric maps $M \to \mathbb{R}^3$ for a polyhedral metric M, so that one could carry out this enumeration and identify the one map that gives a convex polyhedron. In 2005, Fedorchuk and Pak [6] showed an exponential upper bound on the number of such maps. An exponential lower bound is easy to find, so this algorithm takes time exponential in n and is therefore unsatisfactory.

Recent work by Bobenko and Izmestiev [2] produced a new proof of Alexandrov's Theorem, describing a certain ordinary differential equation (ODE) and initial conditions whose solution contains sufficient information to construct the embedding by elementary geometry. This work was accompanied by a computer implementation of the ODE [14], which empirically produces accurate approximations of embeddings of metrics on which it is tested.

In this work, we describe an algorithm based on the Bobenko-Izmestiev ODE, and prove a pseudopolynomial bound on its running time. Specifically, call an embedding of M *ε-accurate* if the metric is distorted by at most a factor $1+\varepsilon$, and *ε-convex* if each dihedral angle is at most $\pi+\varepsilon$. For concreteness, M may be represented by a list of triangles with side lengths and the names of adjacent triangles. Then we show the following theorem:

Theorem 1. *Given a polyhedral metric M with n vertices, ratio S between the largest and smallest distance between vertices, and defect (discrete Gaussian curvature) between ε_1 and $2\pi - \varepsilon_8$ at each vertex, an ε_6-accurate ε_9-convex embedding of M can be found in time $O\left(n^{913/2}S^{831}/(\varepsilon^{121}\varepsilon_1^{445}\varepsilon_8^{616})\right)$ where $\varepsilon = \min(\varepsilon_6/nS, \varepsilon_9\varepsilon_1^2/nS^6)$.*

The exponents in the time bound of Theorem 1 are remarkably large. Thankfully, no evidence suggests our algorithm actually takes as long to run as the bound allows. On the contrary, our analysis relies on bounding approximately a dozen geometric quantities, and to keep the analysis tractable we use the simplest bound whenever available. The algorithm's actual performance is governed by the actual values of these quantities, and therefore by whatever sharper bounds could be proven by a stingier analysis.

To describe our approach, consider an embedding of the metric M as a convex polyhedron in $\mathbb{R}^3$, and choose an arbitrary origin O in the surface's interior. Then it is not hard to see that the n distances $r_i = \overline{Ov_i}$ from the origin to the vertices v_i, together with M and the combinatorial data describing which polygons on M are faces of the polyhedron, suffice to reconstruct the embedding: the tetrahedron formed by O and each triangle is rigid in $\mathbb{R}^3$, and we have no choice in how to glue them to each other. In Lemma 1 below, we show that in fact the radii alone suffice to reconstruct the embedding, to do so efficiently, and to do so even with radii of finite precision.

Therefore in order to compute the unique embedding of M that Alexandrov's Theorem guarantees exists, we compute a set of radii $r = \{r_i\}_i$ and derive a triangulation T. The exact radii satisfy three conditions:

1. the radii r determine nondegenerate tetrahedra from O to each face of T;
2. with these tetrahedra, the dihedral angles at each exterior edge total at most π; and
3. with these tetrahedra, the dihedral angles about each radius sum to 2π.

In our computation, we begin with a set of large initial radii $r_i = R$ satisfying Conditions 1 and 2, and write $\kappa = \{\kappa_i\}_i$ for the differences by which Condition 3 fails about each radius. We then iteratively adjust the radii to bring κ near zero and satisfy Condition 3 approximately, maintaining Conditions 1 and 2 throughout.

The computation takes the following form. We describe the Jacobian $\left(\frac{\partial \kappa_i}{\partial r_j}\right)_{ij}$, showing that it can be efficiently computed and that its inverse is pseudopolynomially bounded. We show further that the Hessian $\left(\frac{\partial \kappa_i}{\partial r_j \partial r_k}\right)_{ijk}$ is also pseudopolynomially bounded. It follows that a change in r in the direction of smaller κ as described by the Jacobian, with some step size only pseudopolynomially small, makes progress in reducing $|\kappa|$. The step size can be chosen online by doubling and halving, so it follows that we can take steps of the appropriate size, pseudopolynomial in number, and obtain an r that zeroes κ to the desired precision in pseudopolynomial total time. Theorem 1 follows.

The construction of [2] is an ODE in the same n variables r_i, with a similar starting point and with the derivative of r driven similarly by a desired path for κ. Their proof differs in that it need only show existence, not a bound, for the Jacobian's inverse, in order to invoke the inverse function theorem. Similarly, while we must show a pseudopolynomial lower bound (Lemma 11) on the altitudes of the tetrahedra during our computation, the prior work shows only that these altitudes remain positive. In general our computation requires that the known open conditions—this quantity is positive, that map is nondegenerate—be replaced by stronger compact conditions—this quantity is lower-bounded, that map's inverse is bounded. We model our proofs of these strengthenings on the proofs in [2] of the simpler open conditions, and we directly employ several other results from that paper where possible.

The remainder of this paper supplies the details of the proof of Theorem 1. We give background in Section 2, and detail the main argument in Section 3. We bound the Jacobian in Section 4 and the Hessian in Section 5. Finally, some lemmas are deferred to Section 6 for clarity.

2 Background and Notation

In this section we define our major geometric objects and give the basic facts about them. We also define some parameters describing our central object that we will need to keep bounded throughout the computation.

2.1 Geometric notions

Central to our argument are two dual classes of geometric structures introduced by Bobenko and Izmestiev in [2] under the names of "generalized convex polytope" and "generalized convex polyhedron". Because in other usages the distinction between "polyhedron" and "polytope" is that a polyhedron is a three-dimensional polytope, and because both of these objects are three-dimensional, we will refer to these objects as "generalized convex polyhedra" and "generalized convex dual polyhedra" respectively to avoid confusion.

First, we define the objects that our main theorem is about.

Definition 1. *A metric M homeomorphic to the sphere is a* polyhedral metric *if each $x \in M$ has an open neighborhood isometric either to a subset of $\mathbb{R}^2$ or to a cone of angle less than 2π with x mapped to the apex, and if only finitely many x, called the* vertices *$V(M) = \{v_i\}_i$ of M, fall into the latter case.*

The defect *δ_i at a vertex $v_i \in V(M)$ is the difference between 2π and the total angle at the vertex, which is positive by the definition of a vertex.*

An embedding *of M is a piecewise linear map $f : M \to \mathbb{R}^3$. An embedding f is ε-*accurate *if it distorts the metric M by at most $1+\varepsilon$, and ε-*convex *if $f(M)$ is a polyhedron and each dihedral angle in $f(M)$ is at most $\pi + \varepsilon$.*

Definition 2. *In a tetrahedron $ABCD$, write $\angle CABD$ for the dihedral angle along edge AB.*

Definition 3. *A* triangulation *of a polyhedral metric M is a decomposition into Euclidean triangles whose vertex set is $V(M)$. Its vertices are denoted by $V(T) = V(M)$, its edges by $E(T)$, and its faces by $F(T)$.*

A radius assignment *on a polyhedral metric M is a map $r : V(M) \to \mathbb{R}_+$. For brevity we write r_i for $r(v_i)$.*

A generalized convex polyhedron *is a gluing of metric tetrahedra with a common apex O. The generalized convex polyhedron $P = (M, T, r)$ is determined by the polyhedral metric M and triangulation T giving its bases and the radius assignment r for the side lengths.*

Write $\kappa_i \triangleq 2\pi - \sum_{jk} \angle v_j O v_i v_k$ for the curvature about Ov_i, and $\phi_{ij} \triangleq \angle v_i O v_j$ for the angle between vertices v_i, v_j seen from the apex.

Our algorithm, following the construction in [2], will choose a radius assignment for the M in question and iteratively adjust it until the associated generalized convex polyhedron P fits nearly isometrically in $\mathbb{R}^3$. The resulting radii will give an ε-accurate ε-convex embedding of M into $\mathbb{R}^3$.

In the argument we will require several geometric objects related to generalized convex polyhedra.

Definition 4. *A* Euclidean simplicial complex *is a metric space on a simplicial complex where the metric restricted to each cell is Euclidean.*

A generalized convex polygon *is a Euclidean simplicial 2-complex homeomorphic to a disk, where all triangles have a common vertex V, the total angle at V is no more than 2π, and the total angle at each other vertex is no more than π.*

Given a generalized convex polyhedron $P = (M, T, r)$, the corresponding generalized convex dual polyhedron $D(P)$ *is a certain Euclidean simplicial 3-complex. Let O be a vertex called the* apex, *A_i a vertex with $OA_i = h_i \triangleq 1/r_i$ for each i.*

For each edge $v_iv_j \in E(T)$ bounding triangles $v_iv_jv_k$ and $v_jv_iv_l$, construct two simplices $OA_iA_{jil}A_{ijk}$, $OA_jA_{ijk}A_{jil}$ in $D(P)$ as follows. Embed the two tetrahedra $Ov_iv_jv_k, Ov_jv_iv_l$ in $\mathbb{R}^3$. For each $i' \in \{i, j, k, l\}$, place $A_{i'}$ along ray $Ov_{i'}$ at distance $h_{i'}$, and draw a perpendicular plane $P_{i'}$ through the ray at $A_{i'}$. Let A_{ijk}, A_{jil} be the intersection of the planes P_i, P_j, P_k and P_j, P_i, P_l respectively.

Now identify the vertices $A_{ijk}, A_{jki}, A_{kij}$ for each triangle $v_iv_jv_k \in F(T)$ to produce the Euclidean simplicial 3-complex $D(P)$. Since the six simplices produced about each of these vertices A_{ijk} are all defined by the same three planes P_i, P_j, P_k with the same relative configuration in $\mathbb{R}^3$, the total dihedral angle about each OA_{ijk} is 2π. On the other hand, the total dihedral angle about OA_i is $2\pi - \kappa_i$, and the face about A_i is a generalized convex polygon of defect κ_i.

The Jacobian bound in Section 4 makes use of certain multilinear forms described in [2] and in the full paper [8].

Definition 5. *The* dual volume $\text{vol}(h)$ *is the volume of the generalized convex dual polyhedron $D(P)$, a cubic form in the dual altitudes h. The* mixed volume $\text{vol}(\cdot,\cdot,\cdot)$ *is the associated symmetric trilinear form.*

Let E_i be the area of the face around A_i in $D(P)$, a quadratic form in the altitudes within this face. The ith mixed area $E_i(\cdot,\cdot)$ *is the associated symmetric bilinear form.*

Let π_i be the linear map $\pi_i(h)_j \triangleq \frac{h_j - h_i \cos\phi_{ij}}{\sin\phi_{ij}}$. so that $\pi_i(h) = g(i)$. Then define $F_i(a, b) \triangleq E_i(\pi_i(a), \pi_i(b))$ so that $F_i(h, h)$ is the area of face i.

By elementary geometry $\text{vol}(h, h, h) = \frac{1}{3}\sum_i h_i F_i(h, h)$, so that by a simple computation $\text{vol}(a, b, c) = \frac{1}{3}\sum_i a_i F_i(b, c)$.

2.2 Weighted Delaunay triangulations

The triangulations we require at each step of the computation are the weighted Delaunay triangulations used in the construction of [2]. We give a simpler definition inspired by Definition 14 of [7].

Definition 6. *The* power $\pi_v(p)$ *of a point p against a vertex v in a polyhedral metric M with a radius assigment r is $pv^2 - r(v)^2$.*

The center $C(v_iv_jv_k)$ *of a triangle $v_iv_jv_k \in T(M)$ when embedded in $\mathbb{R}^2$ is the unique point p such that $\pi_{v_i}(p) = \pi_{v_j}(p) = \pi_{v_k}(p)$, which exists by the radical axis theorem from classical geometry. The quantity $\pi_{v_i}(p) = \pi(v_iv_jv_k)$ is the* power *of the triangle.*

A triangulation T of a polyhedral metric M with radius assignment r is locally convex *at edge v_iv_j with neighboring triangles $v_iv_jv_k, v_jv_iv_l$ if*

$\pi_{v_l}(C(v_iv_jv_k)) > \pi_{v_l}(v_k)$ and $\pi_{v_k}(C(v_jv_iv_l)) > \pi_{v_k}(v_l)$ when $v_iv_jv_k, v_jv_iv_l$ are embedded together in $\mathbb{R}^2$.

A weighted Delaunay triangulation *for a radius assignment r on a polyhedral metric M is a triangulation T that is locally convex at every edge.*

A weighted Delaunay triangulation can be computed in time $O(n^2 \log n)$ by a simple modification of the *continuous Dijkstra algorithm* of [10]. The original analysis of this algorithm assumes that each edge of the input triangulation is a shortest path. In the full paper [8] we show that the same algorithm works in time $O(S\varepsilon_8^{-1}n^2 \log n)$ in the general case. Therefore we perform the general computation once at the outset, and use the resulting triangulation as the basis for subsequent runs of the continuous Dijkstra algorithm in time $O(n^2 \log n)$ each.

The radius assignment r and triangulation T admit a tetrahedron $Ov_iv_jv_k$ just if the power of $v_iv_jv_k$ is negative, and the squared altitude of O in this tetrahedron is $-\pi(v_iv_jv_k)$. The edge v_iv_j is convex when the two neighboring tetrahedra are embedded in $\mathbb{R}^3$ just if it is locally convex in the triangulation as in Definition 6. A weighted Delaunay triangulation with negative powers therefore gives a valid generalized convex polyhedron if the curvatures κ_i are positive. For each new radius assignment r in the computation of Section 3 we therefore compute the weighted Delaunay triangulation and proceed with the resulting generalized convex polyhedron, in which Lemma 11 guarantees a positive altitude and the choices in the computation guarantee positive curvatures.

2.3 Notation for bounds

Definition 7. *Let the following bounds be observed:*

1. *n is the number of vertices on M.*
2. *$\varepsilon_1 \stackrel{\Delta}{=} \min_i \delta_i$ is the minimum defect.*
3. *$\varepsilon_2 \stackrel{\Delta}{=} \min_i(\delta_i - \kappa_i)$ is the minimum defect-curvature gap.*
4. *$\varepsilon_3 \stackrel{\Delta}{=} \min_{ij \in E(T)} \phi_{ij}$ is the minimum angle between radii.*
5. *$\varepsilon_4 \stackrel{\Delta}{=} \max_i \kappa_i$ is the maximum curvature.*
6. *$\varepsilon_5 \stackrel{\Delta}{=} \min_{v_iv_jv_k \in F(T)} \angle v_iv_jv_k$ is the smallest angle in the triangulation. Observe that obtuse angles are also bounded: $\angle v_iv_jv_k < \pi - \angle v_jv_iv_k \leq \pi - \varepsilon_5$.*
7. *ε_6 is used for the desired accuracy in embedding M.*
8. *$\varepsilon_7 \stackrel{\Delta}{=} (\max_i \frac{\kappa_i}{\delta_i})/(\min_i \frac{\kappa_i}{\delta_i}) - 1$ is the extent to which the ratio among the κ_i varies from that among the δ_i. We will keep $\varepsilon_7 < \varepsilon_8/4\pi$ throughout.*
9. *$\varepsilon_8 \stackrel{\Delta}{=} \min_i(2\pi - \delta_i)$ is the minimum angle around a vertex.*
10. *ε_9 is used for the desired approximation to convexity in embedding M.*
11. *D is the diameter of M.*
12. *L is the maximum length of any edge in the input triangulation.*
13. *ℓ is the shortest distance v_iv_j between vertices.*
14. *$S \stackrel{\Delta}{=} \max(D, L)/\ell$ is the maximum ratio of distances.*

15. $d_0 \triangleq \min_{p \in M} Op$ *is the minimum height of the apex off of any point on* M.
16. $d_1 \triangleq \min_{v_i v_j \in E(T)} d(O, v_i v_j)$ *is the minimum distance to any edge of* T.
17. $d_2 \triangleq \min_i r_i$ *is the minimum distance from the apex to any vertex of* M.
18. $H \triangleq 1/d_0$*; the name is justified by* $h_i = 1/r_i \leq 1/d_0$.
19. $R \triangleq \max_i r_i$*, so* $1/H \leq r_i \leq R$ *for all* i.
20. $T \triangleq HR$ *is the maximum ratio of radii.*

Of these bounds, $n, \varepsilon_1, \varepsilon_8$, and S are fundamental to the given metric M or the form in which it is presented as input, and D, L, and ℓ are dimensionful parameters of the same metric input. The values ε_6 and ε_9 define the objective to be achieved, and our computation will drive ε_4 toward zero while maintaining ε_2 large and ε_7 small. In Section 6 we bound the remaining parameters $\varepsilon_3, \varepsilon_5, R, d_0, d_1$, and d_2 in terms of these.

Definition 8. *Let* $\mathbf{J}$ *denote the Jacobian* $\left(\frac{\partial \kappa_i}{\partial r_j}\right)_{ij}$*, and* $\mathbf{H}$ *the Hessian* $\left(\frac{\partial \kappa_i}{\partial r_j \partial r_k}\right)_{ijk}$.

3 Main Theorem

In this section, we prove our main theorem using the results proved in the remaining sections. The algorithm of Theorem 1 obtains an approximate embedding of the polyhedral metric M in $\mathbb{R}^3$. Its main subroutine is described by the following theorem:

Theorem 2. *Given a polyhedral metric* M *with* n *vertices, ratio* S *(the* spread*) between the diameter and the smallest distance between vertices, and defect at least* ε_1 *and at most* $2\pi - \varepsilon_8$ *at each vertex, a radius assignment* r *for* M *with maximum curvature at most* ε *can be found in time* $O\left(n^{913/2} S^{831} / (\varepsilon^{121} \varepsilon_1^{445} \varepsilon_8^{616})\right)$.

Proof. Let a *good* assignment be a radius assignment r that satisfies two bounds: $\varepsilon_7 < \varepsilon_8/4\pi$ so that Lemmas 9–11 apply and r therefore by the discussion in Subsection 2.2 produces a valid generalized convex polyhedron for M, and $\varepsilon_2 = \Omega(\varepsilon_1^2 \varepsilon_8^3 / n^2 S^2)$ on which our other bounds rely. By Lemma 6, there exists a good assignment r^0. We will iteratively adjust r^0 through a sequence r^t of good assignments to arrive at an assignment r^N with maximum curvature $\varepsilon_4^N < \varepsilon$ as required. At each step we recompute T as a weighted Delaunay triangulation according to Subsection 2.2.

Given a good assignment $r = r^n$, we will compute another good assignment $r' = r^{n+1}$ with $\varepsilon_4 - \varepsilon_4' = \Omega\left(\varepsilon_1^{445} \varepsilon_4^{121} \varepsilon_8^{616} / (n^{907/2} S^{831})\right)$. It follows that from r^0 we can arrive at a satisfactory r^N with $N = O\left((n^{907/2} S^{831}) / (\varepsilon^{121} \varepsilon_1^{445} \varepsilon_8^{616})\right)$.

To do this, let $\mathbf{J}$ be the Jacobian $(\frac{\partial \kappa_i}{\partial r_j})_{ij}$ and $\mathbf{H}$ the Hessian $\left(\frac{\partial \kappa_i}{\partial r_j \partial r_k}\right)_{ijk}$, evaluated at r. The goodness conditions and the objective are all in terms of κ, so we choose a desired new curvature vector κ^* in κ-space and apply the inverse Jacobian to get a new radius assignment $r' = r + \mathbf{J}^{-1}(\kappa^* - \kappa)$ in r-space. The actual new curvature vector κ' differs from κ^* by an error at most $\frac{1}{2}|\mathbf{H}||r'-r|^2 \leq \left(\frac{1}{2}|\mathbf{H}||\mathbf{J}^{-1}|^2\right)|\kappa^* - \kappa|^2$, quadratic in the desired change in curvatures with a coefficient

$$C \triangleq \frac{1}{2}|\mathbf{H}||\mathbf{J}^{-1}|^2 = O\left(\frac{n^{3/2}S^{14}}{\varepsilon_5^3}\frac{R^{23}}{D^{14}d_0^3d_1^8}\left(\frac{n^{7/2}T^2}{\varepsilon_2\varepsilon_3^3\varepsilon_4}R\right)^2\right) = O\left(\frac{n^{905/2}S^{831}}{\varepsilon_1^{443}\varepsilon_4^{121}\varepsilon_8^{616}}\right)$$

by Theorems 3 and 4 and Lemmas 7, 6, 11, and 8.

Therefore pick a step size p, and choose κ^* according to $\kappa_i^* - \kappa_i = -p\kappa_i - p\left(\kappa_i - \delta_i \min_j \frac{\kappa_j}{\delta_j}\right)$. The first term diminishes all the curvatures together to reduce ε_4, and the second rebalances them to keep the ratios $\frac{\kappa_j}{\delta_j}$ nearly equal so that ε_7 remains small. In the full paper [8] we show that the resulting actual curvatures κ' make r' a good assignment and put $\varepsilon_4' \leq \varepsilon_4 - p\varepsilon_4/2$, so long as

$$p \leq \varepsilon_1^2/64\pi^2 n\varepsilon_4 C. \tag{1}$$

This produces a good radius assignment r' in which ε_4 has declined by at least

$$\frac{p\varepsilon_4}{2} = \frac{\varepsilon_1^2}{128\pi^2 nC} = \Omega\left(\frac{\varepsilon_1^{445}\varepsilon_4^{121}\varepsilon_8^{616}}{n^{907/2}S^{831}}\right)$$

as required.

As a simplification, we need not compute p exactly according to (1). Rather, we choose the step size p^t at each step, trying first p^{t-1} (with p^0 an arbitrary constant) and computing the actual curvature error $|\kappa' - \kappa^*|$. If the error exceeds its maximum acceptable value $p\varepsilon_1^2\varepsilon_4/16\pi^2$ then we halve p^t and try step t again, and if it falls below half this value then we double p^t for the next round. Since we double at most once per step and halve at most once per doubling plus a logarithmic number of times to reach an acceptable p, this doubling and halving costs only a constant factor. Even more important than the resulting simplification of the algorithm, this technique holds out the hope of actual performance exceeding the proven bounds.

Now each of the N iterations of the computation go as follows. Compute the weighted Delaunay triangulation T^t for r^t in time $O(n^2 \log n)$ as described in Subsection 2.2. Compute the Jacobian $\mathbf{J}^t$ in time $O(n^2)$ using formulas (14, 15) in [2]. Choose a step size p^t, possibly adjusting it, as discussed above. Finally, take the resulting r' as r^{t+1} and continue. The computation of κ^* to check p^t runs in linear time, and that of r' in time $O(n^\omega)$ where $\omega < 3$ is the time exponent of matrix multiplication. Each iteration therefore costs time $O(n^3)$, and the whole computation costs time $O(n^3 N)$ as claimed. □

Now with our radius assignment r for M and the resulting generalized convex polyhedron P with curvatures all near zero, it remains to approximately embed P and therefore M in $\mathbb{R}^3$. To begin, we observe that this is easy to do given exact values for r and in a model with exact computation: after triangulating, P is made up of rigid tetrahedra and we embed one tetrahedron arbitrarily, then embed each neighboring tetrahedron in turn.

In a realistic model, we compute only with bounded precision, and in any case Theorem 2 gives us only curvatures near zero, not equal to zero. Lemma 1 produces an embedding in this case, settling for less than exact isometry and exact convexity.

Lemma 1. *There is an algorithm that, given a radius assignment r for which the corresponding curvatures κ_i are all less than $\varepsilon = O\left(\min(\varepsilon_6/nS, \varepsilon_9\varepsilon_1^2/nS^6)\right)$ for some constant factor, produces explicitly by vertex coordinates in time $O(n^2 \log n)$ an ε_6-accurate ε_9-convex embedding of M.*

Proof (sketch). As in the exact case, triangulate M, embed one tetrahedron arbitrarily, and then embed its neighbors successively. The positive curvature will force gaps between the tetrahedra. Then replace the several copies of each vertex by their centroid, so that the tetrahedra are distorted but leave no gaps. This is the desired embedding. The proofs of ε_6-accuracy and ε_9-convexity are straightforward and left to the full paper [8].

A weighted Delaunay triangulation takes time $O(n^2 \log n)$ as discussed in Subsection 2.2, and the remaining steps take time $O(n)$. □

We now have all the pieces to prove our main theorem.

Proof (Theorem 1). Let $\varepsilon \stackrel{\Delta}{=} O\left(\min(\varepsilon_6/nS, \varepsilon_9\varepsilon_1^2/nS^6)\right)$, and apply the algorithm of Theorem 2 to obtain in time $O\left(n^{913/2}S^{831}/(\varepsilon^{121}\varepsilon_1^{445}\varepsilon_8^{616})\right)$ a radius assignment r for M with maximum curvature $\varepsilon_4 \le \varepsilon$.

Now apply the algorithm of Lemma 1 to obtain in time $O(n^2 \log n)$ the desired embedding and complete the computation. □

4 Bounding the Jacobian

Theorem 3. *The Jacobian $\mathbf{J} = \left(\frac{\partial \kappa_i}{\partial r_j}\right)_{ij}$ has inverse pseudopolynomially bounded by $|\mathbf{J}^{-1}| = O\left(\frac{n^{7/2}T^2}{\varepsilon_2\varepsilon_3^3\varepsilon_4}R\right)$.*

Proof. Our argument parallels that of Corollary 2 in [2], which concludes that the same Jacobian is nondegenerate. Theorem 4 of [2] shows that this Jacobian equals the Hessian of the volume of the dual $D(P)$. The meat of the corollary's proof is in Theorem 5 of [2], which begins by equating this Hessian to the bilinear form $6\,\mathrm{vol}(h,\cdot,\cdot)$ derived from the mixed volume we defined in Definition 5. So we have to bound the inverse of this bilinear form.

To do this it suffices to show that the form $\mathrm{vol}(h,x,\cdot)$ has norm at least $\Omega\left(\frac{\varepsilon_2\varepsilon_3^3\varepsilon_4}{n^{7/2}T^2}\frac{|x|}{R}\right)$ for all vectors x. Equivalently, suppose some x has $|\mathrm{vol}(h,x,z)| \le |z|$ for all z; we show $|x| = O\left(\frac{n^{7/2}T^2}{\varepsilon_2\varepsilon_3^3\varepsilon_4}R\right)$.

To do this we follow the proof in Theorem 5 of [2] that the same form $\mathrm{vol}(h,x,\cdot)$ is nonzero for x nonzero. Throughout the argument we work in terms of the dual $D(P)$.

Recall that for each i, $\pi_i x$ is defined as the vector $\{x_{ij}\}_j$. It suffices to show that for all i

$$|\pi_i x|_2^2 = O\left(\frac{n^3T^3}{\varepsilon_2^2\varepsilon_3\varepsilon_4}R^2 + \frac{n^2T^2}{\varepsilon_2\varepsilon_3\varepsilon_4}R|x|_1\right)$$

since then by Lemma 2

$$|x|_2^2 \leq \frac{4n}{\varepsilon_3^2}\max_i |\pi_i x|_2^2 = O\left(\frac{n^4T^3}{\varepsilon_2^2\varepsilon_3^3\varepsilon_4}R^2 + \frac{n^3T^2}{\varepsilon_2\varepsilon_3^3\varepsilon_4}R|x|_1\right),$$

and since $|x|_1 \leq \sqrt{n}|x|_2$ and $X^2 \leq a + bX$ implies $X \leq \sqrt{a} + b$, $|x|_2 = O\left(\frac{n^{7/2}T^2}{\varepsilon_2\varepsilon_3^3\varepsilon_4}R\right)$. Therefore fix an arbitrary i, let $g = \pi_i h$ and $y = \pi_i x$, and we proceed to bound $|y|_2$.

We break the space on which E_i acts into the 1-dimensional positive eigenspace of E_i and its $(k-1)$-dimensional negative eigenspace, since by Lemma 3.4 of [2] the signature of E_i is $(1, k-1)$, where k is the number of neighbors of v_i. Write λ_+ for the positive eigenvalue and $-E_i^-$ for the restriction to the negative eigenspace so that E_i^- is positive definite, and decompose $g = g_+ + g_-$, $y = y_+ + y_-$ by projection into these subspaces. Then we have

$$\begin{aligned}
G \triangleq E_i(g,g) &= \lambda_+ g_+^2 - E_i^-(g_-,g_-) \triangleq \lambda_+ g_+^2 - G_- \\
E_i(g,y) &= \lambda_+ g_+ y_+ - E_i^-(g_-,y_-) \\
Y \triangleq E_i(y,y) &= \lambda_+ y_+^2 - E_i^-(y_-,y_-) \triangleq \lambda_+ y_+^2 - Y_-
\end{aligned}$$

and our task is to obtain an upper bound on $Y_- = E_i^-(y_-,y_-)$, which will translate through our bound on the eigenvalues of E_i away from zero into the desired bound on $|y|$.

We begin by obtaining bounds on $|E_i(g,y)|$, G_-, G, and Y. Since $|z| \geq |\mathrm{vol}(h,x,z)|$ for all z and $\mathrm{vol}(h,x,z) = \sum_j z_j F_j(h,x)$, we have $|E_i(g,y)| = |F_i(h,x)| \leq 1$. Further, $\det\begin{pmatrix} E_i(g,g) & E_i(y,g) \\ E_i(g,y) & E_i(y,y)\end{pmatrix} < 0$ because E_i has signature $(1,1)$ restricted to the (y,g) plane, so by Lemma 3 $Y = E_i(y,y) < \frac{R^2}{\varepsilon_2}$.

Now by further calculation and the use of Lemma 4, the theorem follows; the details are left to the full paper [8] for brevity. □

Three small lemmas used above follow from the geometry of spherical polygons and of generalized convex dual polyhedra. Their proofs are left to the full paper [8] for brevity.

Lemma 2. $|x|^2 \leq (4n/\varepsilon_3^2)\max_i |\pi_i x|^2$.

Lemma 3. $F_i(h,h) > \varepsilon_2/R^2$.

Lemma 4. *The inverse of the form E_i is bounded by $|E_i^{-1}| = O(n/\varepsilon_4)$.*

5 Bounding the Hessian

In order to control the error in each step of our computation, we need to keep the Jacobian $\mathbf{J}$ along the whole step close to the value it started at, on which the step was based. To do this we bound the Hessian $\mathbf{H}$ when the triangulation is fixed, and we show that the Jacobian does not change discontinuously when changing radii force a new triangulation.

Each curvature κ_i is of the form $2\pi - \sum_{j,k:v_iv_jv_k\in T} \angle v_jOv_iv_k$, so in analyzing its derivatives we focus on the dihedral angles $\angle v_jOv_iv_k$. When the tetrahedron $Ov_iv_jv_k$ is embedded in $\mathbb{R}^3$, the angle $\angle v_jOv_iv_k$ is determined by elementary geometry as a smooth function of the distances among O, v_i, v_j, v_k. For a given triangulation T this makes κ a smooth function of r. Our first lemma shows that no error is introduced at the transitions where the triangulation $T(r)$ changes.

Lemma 5. *The Jacobian* $\mathbf{J} = \left(\frac{\partial\kappa_i}{\partial r_j}\right)_{ij}$ *is continuous at the boundary between radii corresponding to one triangulation and to another.*

Proof (sketch). The proof, which can be found in the full paper [8], uses elementary geometry to compare the figures determined by two triangulations near a radius assignment on their boundary. $\square$

It now remains to control the change in $\mathbf{J}$ as r changes within any particular triangulation, which we do by bounding the Hessian.

Theorem 4. *The Hessian* $\mathbf{H} = \left(\frac{\partial\kappa_i}{\partial r_j\partial r_k}\right)_{ijk}$ *is bounded in norm by* $O\big(n^{5/2}S^{14}R^{23}/(\varepsilon_5^3 d_0^3 d_1^8 D^{14})\big)$.

Proof. By direct computation and computer algebra. See the full paper [8] for the details. $\square$

6 Intermediate Bounds

Here we bound miscellaneous parameters in the computation in terms of the fundamental parameters $n, S, \varepsilon_1, \varepsilon_8$ and the computation-driving parameter ε_4.

Lemma 6. *Given a polyhedral metric space M, there exists a radius assignment r with curvature skew $\varepsilon_7 < \varepsilon_8/4\pi$, maximum radius $R = O(nD/\varepsilon_1\varepsilon_8)$, and minimum defect-curvature gap $\varepsilon_2 = \Omega(\varepsilon_1^2\varepsilon_8^3/n^2S^2)$.*

Proof (sketch). Take $r_i = R$ for all i, with R sufficiently large. Then each κ_i is nearly equal to δ_i, so that ε_7 is small. For the quantitative bounds and a complete proof, see the full paper [8]. $\square$

Two bounds on angles can be proven by elementary geometry; details are left to the full paper [8] for brevity.

Lemma 7. $\varepsilon_3 > \ell d_1/R^2$.

Lemma 8. $\varepsilon_5 > \varepsilon_2/6S$.

Finally we bound O away from the surface M. The bounds are effective versions of Lemmas 4.8, 4.6, and 4.5 respectively of [2], and the proofs, left for brevity to the full paper [8], are similar but more involved.

Recall that d_2 is the minimum distance from O to any vertex of M, d_1 is the minimum distance to any edge of T, and d is the minimum distance from O to any point of M.

Lemma 9. $d_2 = \Omega\big(D\varepsilon_1\varepsilon_4\varepsilon_5^2\varepsilon_8/(nS^4)\big)$.

Lemma 10. $d_1 = \Omega\big(D\varepsilon_1^2\varepsilon_4^4\varepsilon_5^6\varepsilon_8^2/(n^2S^{10})\big)$.

Lemma 11. $d_0 = \Omega\big(D\varepsilon_1^4\varepsilon_4^9\varepsilon_5^{12}\varepsilon_8^4/(n^4S^{22})\big)$.

Acknowledgments. We thank Jeff Erickson and Joseph Mitchell for helpful discussions about shortest paths on non-shortest-path triangulations.

References

1. A. D. Alexandrov. *Convex Polyhedra*. Springer-Verlag, Berlin, 2005.
2. A. I. Bobenko and I. Izmestiev. Alexandrov's theorem, weighted Delaunay triangulations, and mixed volumes. *Annales de l'Institut Fourier*, 2006, arXiv:math.DG/0609447.
3. A. L. Cauchy. Sur les polygones et les polyèdres, seconde mémoire. *J. École Polytechnique*, XVIe Cahier, Tome IX:113–148, 1813. In OEuvres Complètes, IIe Sèrie, volume 1, pages 26–38, Paris, 1905.
4. E. D. Demaine, M. L. Demaine, A. Lubiw, and J. O'Rourke. Enumerating foldings and unfoldings between polygons and polytopes. *Graphs Comb.*, 18:93–104, 2002.
5. E. D. Demaine and J. O'Rourke. *Geometric Folding Algorithms*. Cambridge University Press, Cambridge, 2007.
6. M. Fedorchuk and I. Pak. Rigidity and polynomial invariants of convex polytopes. *Duke Math. J.*, 129:371–404, 2005.
7. David Glickenstein. Geometric triangulations and discrete Laplacians on manifolds, 2005, arXiv:math/0508188v1.
8. Daniel Kane, Gregory N. Price, and Erik D. Demaine. A pseudopolynomial algorithm for Alexandrov's theorem. 2008, arXiv:0812.5030.
9. A. Lubiw and J. O'Rourke. When can a polygon fold to a polytope? Technical Report 048, Department of Computer Science, Smith College, 1996. Presented at *Am. Math. Soc.* Conf., 5 October 1996.
10. J. S. B. Mitchell, D. M. Mount, and C. H. Papadimitriou. The discrete geodesic problem. *SIAM J. Comput.*, 16:647–668, 1987.
11. I. Sabitov. The volume of a polyhedron as a function of its metric. *Fundam. Prikl. Mat.*, 2(4):1235–1246, 1996.
12. I. Sabitov. The volume of a polyhedron as a function of its metric and algorithmical solution of the main problems in the metric theory of polyhedra. In *International School-Seminar Devoted to the N. V. Efimov's Memory*, pages 64–65, 1996.
13. I. Sabitov. The volume as a metric invariant of polyhedra. *Discrete Comput. Geom.*, 20:405–425, 1998.
14. Stefan Sechelmann. Alexandrov polyhedron editor, 2006, http://www.math.tu-berlin.de/geometrie/ps/software.shtml#AlexandrovPolyhedron.