

Convex optimization on CAT(0) cubical complexes

Ariel Goodwin^{*} Adrian S. Lewis[†] Genaro López-Acedo[‡]
Adriana Nicolae[§]

May 3, 2024

Abstract

We consider geodesically convex optimization problems involving distances to a finite set of points A in a CAT(0) cubical complex. Examples include the minimum enclosing ball problem, the weighted mean and median problems, and the feasibility and projection problems for intersecting balls with centers in A . We propose a decomposition approach relying on standard Euclidean cutting plane algorithms. The cutting planes are readily derivable from efficient algorithms for computing geodesics in the complex.

Key words: convex, cutting plane, geodesic, CAT(0), cubical complex

AMS Subject Classification: 90C48, 52A41, 57Z25, 65K05

1 Introduction: optimization in Hadamard space

Metric spaces in which points can always be joined by geodesics — isometric images of real intervals — support a wide variety of interesting convex optimization problems [4,31]. Convex optimization is best understood in *Hadamard spaces*: complete geodesic spaces satisfying the *CAT(0) inequality*, which requires that squared-distance functions to given points are all 1-strongly convex. Manifolds comprise the most familiar examples, including Hilbert space, hyperbolic space, and the space of positive-definite matrices with its affine-invariant metric.

^{*}CAM, Cornell University, Ithaca, NY. [arielgoodwin.github.io](https://github.com/arielgoodwin)

[†]ORIE, Cornell University, Ithaca, NY. people.orie.cornell.edu/aslewis

Research supported in part by National Science Foundation Grant DMS-2006990.

[‡]Department of Mathematical Analysis – IMUS, University of Seville, 41012 Seville, Spain
glopez@us.es

[§]Department of Mathematics, Babeş-Bolyai University, 400084, Cluj-Napoca, Romania
anicolae@math.ubbcluj.ro

From a computational perspective, in the special case of manifolds, gradient-based methods are available [8]. On the other hand, even with no differentiable structure, some Hadamard spaces allow efficient computation of geodesics. Particularly interesting are phylogenetic tree spaces [7] and more general CAT(0) cubical complexes [2]. In these spaces, geodesics can be computed in polynomial time [16, 26]. Along with phylogenetic models, applications include reconfigurable systems in robotics [1, 12].

Unfortunately, optimization algorithms in the setting of a general Hadamard space are scarce and slow. Methods based on alternating projections are sometimes available [5], but other general-purpose algorithms typically rely on some predetermined sequence of step sizes. The quintessential example is the problem of computing the mean of a finite set A in the geodesic space (\mathbf{X}, d) , which is the unique minimizer of the function

$$(1.1) \quad m_A(x) = \sum_{a \in A} d^2(a, x) \quad (x \in \mathbf{X}).$$

To compute the mean, the only known general methods iteratively update the n th iterate x_n to a point on the geodesic between x_n and some point $a_n \in A$, chosen randomly [28] or via some deterministic strategy [6, 17, 22]: a standard example is

$$x_{n+1} = \frac{n}{n+1}x_n + \frac{1}{n+1}a_n.$$

To compute the mean of the set $\{-1, +1\} \subset \mathbf{R}$ by this method, for example, starting at the point $x_1 = \frac{1}{2}$ and using $a_n = (-1)^n$, results in the slowly converging sequence $x_n = \frac{1}{2n}(-1)^{n+1}$. Computation in practice confirms this slow convergence: see [24] for experiments in the manifold of positive definite matrices, and [10] for experiments in phylogenetic tree space.

In this work, we focus on the particular case of geodesically convex optimization on a CAT(0) cubical complex. This setting has several features suggesting faster algorithmic possibilities. First, as we have noted, computing geodesics is tractable. Secondly, the underlying space, by definition, decomposes into Euclidean cubes. Lastly, restricted to each cube, geodesically convex optimization problems reduce to standard Euclidean convex optimization. Using these ingredients, we suggest a new cutting plane approach. Our central technique uses geodesics in CAT(0) cubical complexes to derive Euclidean subgradients of distance functions restricted to individual cells in the complex. We illustrate the new algorithm on a simple computational example.

2 CAT(0) cubical complexes

We begin with some standard definitions [9]. Let (\mathbf{X}, d) be a metric space. A *geodesic path* is a distance-preserving mapping $\gamma : [0, l] \subseteq \mathbf{R} \rightarrow \mathbf{X}$. The image

$\gamma([0, l])$ is called a *geodesic segment*. We say that \mathbf{X} is a (*uniquely*) *geodesic space* if every two points in \mathbf{X} are joined by a (unique) geodesic segment. A function $f: \mathbf{X} \rightarrow \mathbf{R}$ is *convex* when its composition with every geodesic path is convex.

Given three points x, y, z in a geodesic space, a *geodesic triangle* $\Delta = \Delta(x, y, z)$ is the union of three geodesic segments (its sides) joining each pair of points. A *comparison triangle* for Δ is a triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ in \mathbf{R}^2 that has side lengths equal to those of Δ . A geodesic metric space is called *CAT(0)* if in every geodesic triangle, distances between points on its sides are no larger than corresponding distances in a comparison triangle. CAT(0) spaces are uniquely geodesic. A *Hadamard space* is a complete CAT(0) space. For any nonempty closed convex subset C of a Hadamard space \mathbf{X} , every point in \mathbf{X} has a unique nearest point in C .

Let $\gamma: [0, l] \rightarrow \mathbf{X}$ and $\eta: [0, r] \rightarrow \mathbf{X}$ be two nonconstant geodesics paths in a CAT(0) space issuing at the same point $x = \gamma(0) = \eta(0)$. For $t \in [0, l]$ and $s \in [0, r]$, denote $\gamma_t = \gamma(t)$ and $\eta_s = \eta(s)$, and let $\Delta(\bar{\gamma}_t, \bar{x}, \bar{\eta}_s)$ be a comparison triangle for $\Delta(\gamma_t, x, \eta_s)$. Then the angle $\angle \bar{\gamma}_t \bar{x} \bar{\eta}_s$ is a nondecreasing function of both t and s , and the *Alexandrov angle* between γ and η is defined by

$$\angle(\gamma, \eta) = \angle \gamma_l x \eta_r = \lim_{t, s \searrow 0} \angle \bar{\gamma}_t \bar{x} \bar{\eta}_s.$$

A *polyhedral cell* C is a geodesic metric space isometric to the convex hull \hat{C} of finitely many points in a Euclidean space \mathbf{R}^n : we usually identify C and \hat{C} . By a *face* of C , we mean a nonempty set that is either C itself, or is the intersection of C with a hyperplane H such that C belongs to one of the closed half-spaces determined by H . The *dimension* of a face is the dimension of the intersection of all affine subspaces containing it. The 0-dimensional faces of a cell are called its *vertices*, and the 1-dimensional faces are called its *edges*. A *polyhedral complex* is a set of polyhedral cells of various dimensions such that the face of any cell is also a cell of the complex, and the intersection of any two cells is either empty or a face of both. The complex is *finite* if it consists of finitely many cells, and is *cubical* if each n -dimensional cell is isometric to the unit cube $[0, 1]^n$.

Given two points x, y in a polyhedral complex \mathbf{X} , the distance $d(x, y)$ is the infimum of lengths of piecewise geodesic paths joining x to y . A *piecewise geodesic path* from x to y is an ordered sets of points $x_0 = x, x_1, \dots, x_k = y$ in \mathbf{X} such that for every $i \in \{1, \dots, k\}$, there exists a cell C_i with $x_{i-1}, x_i \in C_i$; its length is $\sum_{i=1}^k |x_{i-1} - x_i|$. If \mathbf{X} is connected and finite, then (\mathbf{X}, d) is a complete geodesic metric space.

In what follows, we will, in general, consider finite cubical complexes. As proved by Gromov [14], a cubical complex is CAT(0) if and only if it is simply connected and satisfies the following “link” condition at each vertex. Consider the edges containing that vertex. Any set of k distinct such edges, each pair of which is contained in a common 2-dimensional cell, must consist of edges all contained in a common k -dimensional cell. For more discussion, see [16].

Example 2.1 (A simple CAT(0) cubical complex) Consider the space

$$\mathbf{X} = \{x \in [-1, 1]^2 : x_1 \leq 0 \text{ or } x_2 \leq 0\},$$

with the distance induced by the Euclidean metric: in other words, the distance between points is the Euclidean length of the shortest path between them in \mathbf{X} . This space is a finite CAT(0) cubical complex, consisting of the three 2-dimensional cells

$$P_1 = [-1, 0] \times [0, 1], \quad P_2 = [-1, 0] \times [-1, 0], \quad P_3 = [0, 1] \times [0, -1].$$

More generally than the example above, any simply connected subcomplex of \mathbf{Z}^2 , the integer lattice cubing of \mathbf{R}^2 , is CAT(0), because no three edges containing a vertex can, pairwise, be edges of common squares, so the link condition holds. On the other hand, consider the cubical complex formed from the cube $[0, 1]^3$ in \mathbf{R}^3 by taking just the three faces containing zero. This simply connected cubical complex is not CAT(0) because it fails the link condition: the edges connecting zero with the three standard unit vectors are pairwise contained in common squares, but no cell contains all three.

3 Decomposition

Before describing a decomposition approach to convex optimization on cubical complexes, we first illustrate by considering the mean of three points in the simple space of Example 2.1.

Example 3.1 (A simple mean calculation) In the CAT(0) cubical complex \mathbf{X} described in Example 2.1, consider the set A consisting of the standard unit vectors e_1 and e_2 , along with $-e_1$. To compute the mean of A , we must solve the underlying optimization problem (1.1). To do so, we decompose \mathbf{X} into the union of three cells P_1 , P_2 , and P_3 , and solve the restricted problem over each cell in turn:

$$\min_{x \in P_1} \{(1 + |x|)^2 + |x - e_2|^2 + |x + e_1|^2\}$$

$$\min_{x \in P_2} \{|x - e_1|^2 + |x - e_2|^2 + |x + e_1|^2\}$$

$$\min_{x \in P_3} \{|x - e_1|^2 + (1 + |x|)^2 + |x + e_1|^2\}.$$

Worth noting is that the first and third problems are not smooth. The optimal solution of both the second and third problems is the point $(0, 0)$. However, the optimal solution of the first problem is the strictly better point $(-\alpha, \alpha)$, where $\alpha = \frac{1}{6}(2 - \sqrt{2})$. This point is therefore the mean.

This example illustrates a significant feature of mean calculations in cubical complexes or other Hadamard spaces that are not manifolds. Even when we can compute geodesics efficiently, that tool alone does not immediately allow us to recognize whether or not a given point is the mean, let alone compute the mean. For example, along the geodesics between the point $(0, 0)$ and each point in the set A , the objective m_A defined by equation (1.1) is minimized at $(0, 0)$. However, $(0, 0)$ is not the mean: it does not minimize m_A over the whole space \mathbf{X} .

Returning to our general problem, we can minimize a convex function f over a finite cubical complex \mathbf{X} by minimizing f over each of the finitely many cells comprising \mathbf{X} separately. To each of these subproblems we can apply a standard algorithm for Euclidean convex minimization. Rather than exhaustively optimizing over every cube, we can instead consider the following conceptual method, inspired by a somewhat analogous conceptual approach sketched in [25, Algorithm 4.4].

Algorithm 3.2 (Minimize convex f on cubical complex \mathbf{X})

```

input: initial point  $x \in \mathbf{X}$ 
 $\mathcal{P} = \emptyset$                                      % set of optimized cells
for iteration = 1, 2, 3, ... do
     $\mathcal{Q} = \{\text{cells } P \notin \mathcal{P} : x \in P\}$            % unoptimized cells containing  $x$ 
    if  $\mathcal{Q} = \emptyset$  then
        return  $x$                                        %  $x$  optimal
    end if
    choose  $P \in \mathcal{Q}$ 
    choose  $x_P$  minimizing  $f$  over  $P$                    % solve new subproblem
    if  $f(x_P) < f(x)$  then
         $x = x_P$                                        % best point so far
    end if
     $\mathcal{P} = \mathcal{P} \cup \{\text{faces of } P\}$                  % update set of optimized cells
end for

```

While unnecessary formally, we would naturally always choose a *maximal* cell P , meaning that no strictly larger cell contains x .

Proposition 3.3 (Termination) *For any finite cubical complex \mathbf{X} and any convex function $f: \mathbf{X} \rightarrow \mathbf{R}$, Algorithm 3.2 terminates, returning a minimizer of f .*

Proof The procedure terminates, since the space \mathbf{X} is a finite complex. Throughout the procedure, at the current iterate x , the value $f(x)$ never increases. At termination, x therefore minimizes the objective f over every cell containing x . Since f is convex, x therefore minimizes it over the whole space \mathbf{X} . \square

As an example, consider the simple example above, on the space (2.1). In Algorithm 3.2, as soon as we choose the cube P_1 , the procedure terminates with the iterate x equal to the mean, whether or not P_2 or P_3 have already been searched.

Means in metric trees

It is illuminating to consider the behavior of Algorithm 3.2 for computing means for the simplest class of CAT(0) cubical complex: the case when the space \mathbf{X} is a *finite metric tree*. In that case, \mathbf{X} consists of the edges and vertices of a finite connected acyclic graph, where we identify edges with 1-dimensional cells of unit length, intersecting at common 0-dimensional cells — the vertices. The Gromov link condition holds trivially.

We consider the problem of computing the mean of a finite subset A of finite metric tree. We illustrate with the following example [18, Example 1], which is a special case of the “open books” [18] discussed in Appendix A.

Example 3.4 (Stickiness) The *3-spider* is the finite metric tree consisting of three copies of the interval $[0, 1]$ joined at the shared origin 0. Denote the three copies P^1, P^2, P^3 , and consider a set A consisting of three points $a^i \in \text{int } P^i$ satisfying

$$a^i < \sum_{j \neq i} a^j \quad \text{for } i = 1, 2, 3.$$

(One particular example is $a^i = \frac{1}{2}$ for each i .) A quick calculation shows that the mean is the shared origin 0. Unlike in the Euclidean case, the mean is insensitive to small changes in the points a^i : it is *sticky* in the sense of [18].

Consider a general finite metric tree \mathbf{X} , with vertex set V and edge set E . Via a finite computation, we can exactly compute the mean of a finite subset A using Algorithm 3.2 to minimize the function m_A in equation (1.1). The data of the problem, in addition to the graph (V, E) , consists, for each point $a \in A$, of an endpoint $v_a \in V$ of an edge $e_a \in E$ containing a , and the distance δ_a between a and v_a along the edge e_a .

Problems of this kind have been widely studied in the operations research literature, in the context of facility location problems [15]. More typical than the mean problem in that context are the “median” or “minimax” problems, involving the sum or maximum of the distance functions rather than sum of their squares, although [15, Section 3.4] notes an efficient algorithm for the mean problem due to Goldman. Goldman’s 1972 algorithm [13] for the “1-center problem on a tree network”, in the terminology of [29], involves a “trichotomy” at each iteration: after checking an edge the algorithm stops or is confined to one of the two subtrees resulting from deleting that edge.

Algorithm 3.2, it transpires, has the same trichotomy property. At the outset of each iteration, we have a current vertex $v \in V$, and a current set $\mathcal{E} \subset E$ of already optimized edges. For each point $a \in A$, we first find the unique sequence P_a of edges joining v to v_a : its cardinality $|P_a|$ is the distance $d(v, v_a)$. The distance from a to

v is therefore given by

$$d(a, v) = \begin{cases} |P_a| + \delta_a & (e_a \notin P_a) \\ |P_a| - \delta_a & (e_a \in P_a). \end{cases}$$

We next choose a vertex v' neighboring the vertex v and with corresponding edge $e = vv'$ outside the set \mathcal{E} , terminating if there is no such edge. We then minimize m_A over e . Denote the set of those points $a \in A$ for which $e \in P_a$ by \hat{A} . If we identify e with the unit interval $[0, 1]$, where v corresponds to the point 0 and v' corresponds to the point 1, then for any point x on $e = [0, 1]$, we have

$$d(a, x) = \begin{cases} d(a, v) + x & (e \notin P_a) \\ d(a, v) - x & (e \in P_a). \end{cases}$$

We now find the unique point $x \in [0, 1]$ minimizing the strictly convex function

$$\sum_{a \in \hat{A}} (d(a, v) - x)^2 + \sum_{a \notin \hat{A}} (d(a, v) + x)^2.$$

The unconstrained minimizer is

$$\bar{x} = \frac{1}{|A|} \left(\sum_{a \in \hat{A}} d(a, v) - \sum_{a \notin \hat{A}} d(a, v) \right).$$

If $0 < \bar{x} < 1$, then the algorithm terminates: the mean is the point on e at a distance \bar{x} from v . Otherwise, we update \mathcal{E} to include e , update the current iterate to v' if $\bar{x} \geq 1$, and repeat.

Solving the subproblems

In the case of mean computations for metric trees, Algorithm 3.2 involves one-dimensional subproblems with closed-form solutions. In general, however, the subproblems are multivariate, requiring iterative techniques.

The approach outlined in [25, Algorithm 4.4] for computing means in the phylogenetic tree space of [7] relies on a smooth but nonconvex interior-penalty philosophy, the complexity of which is unclear. By contrast, our approach via Algorithm 3.2 generates subproblems that, while nonsmooth, are convex. At each iteration of Algorithm 3.2 we can identify the cell P isometrically with a cube $[0, 1]^n$, for some dimension n , and then apply a linearly convergent Euclidean cutting plane algorithm, efficient in theory and reliable in practice.

One such cutting plane approach for general convex objectives f would be to apply the randomized method of [19], which needs just $\tilde{O}(n^2)$ evaluations of $f(x)$ to approximately solve the minimization problem over the cube $[0, 1]^n$. More precisely, if $\max_P f - \min_P f \leq 1$, then with constant probability the excess $f(x) - \min_P f$

is reduced to $\epsilon > 0$ after no more than $O\left(n^2 \log^{O(1)}\left(\frac{n\sqrt{n}}{\epsilon}\right)\right)$ function evaluations. In essence, the method is subgradient-based: at each iteration, it approximates a subgradient, using $\tilde{O}(n)$ function evaluations.

In this work, however, we are primarily interested in structured functions f composed simply from distance functions: a typical example is the mean objective m_A in equation (1.1). Our central observation is that such objectives support methods based on explicit subgradients, an approach with three potential advantages. First, we arrive at a deterministic rather than randomized algorithm. Secondly, the complexity of the available algorithms is better: the classical ellipsoid method still requires $\tilde{O}(n^2)$ function evaluations, but more recent cutting plane algorithms improve this to $\tilde{O}(n)$ [30]. Lastly, we can experiment with algorithms known to be effective in practice, like proximal or level bundle methods [20].

4 Subgradients of distance functions

In a CAT(0) cubical complex \mathbf{X} , the geodesic between any two points, a and x , is computable in polynomial time [2]. Suppose that x lies in a cell P , and consider the distance function to a , restricted to P . We next show how to use the geodesic to calculate a subgradient of this function at the point x .

We start with a simple tool.

Lemma 4.1 *Let \mathbf{X} be a CAT(0) space and consider three points $a, x, w \in \mathbf{X}$ with $a \neq x$ and $w \neq x$. Then*

$$\frac{d(a, x) - d(a, w)}{d(x, w)} \leq \cos(\angle axw).$$

Proof From the triangle inequality we know

$$|d(a, x) - d(a, w)| \leq d(x, w).$$

We deduce

$$d^2(a, x) - 2d(a, x)d(a, w) + d^2(a, w) \leq d^2(x, w),$$

and hence

$$2d(a, x)(d(a, x) - d(a, w)) \leq d^2(a, x) + d^2(x, w) - d^2(a, w).$$

Consequently we have

$$\frac{d(a, x) - d(a, w)}{d(x, w)} \leq \frac{d(a, x)^2 + d(x, w)^2 - d(a, w)^2}{2d(a, x)d(x, w)} \leq \cos(\angle axw),$$

where the second inequality follows from the law of cosines. \square

Theorem 4.2 *In a CAT(0) cubical complex \mathbf{X} , consider two cells P and Q with a common face F , and points $x \in F$ and $y \in Q$. Denote the nearest point to y in P by z . Then z is also the nearest point to y in F . Furthermore, for any point $w \in P$, if $w \neq x$ and $x \neq z$, then the following angle inequality holds:*

$$(4.3) \quad \cos(\angle yxw) \leq \cos(\angle yxz) \cdot \cos(\angle zwx).$$

Before proving the theorem, it helps to keep in mind a simple example.

Example 4.4 Consider any cubical subcomplex \mathbf{X} of \mathbf{Z}^3 containing the cells

$$P = [0, 1] \times [0, 1] \times [0, 1] \quad \text{and} \quad Q = [-1, 0] \times [0, 1] \times \{0\}.$$

Their common face is the cell

$$F = \{0\} \times [0, 1] \times \{0\}.$$

Consider the two points

$$x = \left(0, \frac{1}{3}, 0\right) \in F \quad \text{and} \quad y = \left(-1, \frac{2}{3}, 0\right) \in Q.$$

The closest point to y in P , in either the Euclidean distance or the distance it induces in \mathbf{X} , is

$$z = \left(0, \frac{2}{3}, 0\right) \in F.$$

At x , the angle between the nontrivial geodesic $[x, y]$ and any nontrivial geodesic $[x, w]$ with $w \in P$ is minimized (either in \mathbf{X} or in \mathbf{R}^n) when $w = z$. In fact, the stronger inequality (4.3) holds.

Proof of Theorem 4.2. The result only involves geodesics with endpoints in P and Q . All such geodesics lie in the interval between P and Q , in the terminology of [2], so without loss of generality we can suppose by [2, Theorem 3.5] that \mathbf{X} is a subcomplex of \mathbf{Z}^n , the integer lattice cubing of \mathbf{R}^n .

Without loss of generality, we can suppose that P and Q are unit cubes in \mathbf{Z}^n that are contained in $[-1, 1]^n$ and both contain 0. Corresponding to any two partitions of the index set $\{1, 2, 3, \dots, n\}$ into disjoint subsets,

$$I_{\leq} \cup I_{=} \cup I_{\geq} \quad \text{and} \quad I^{\leq} \cup I^{\doteq} \cup I^{\geq},$$

we can define such cubes in $[-1, 1]^n$ using the relationship set $\Omega = \{\leq, =, \geq\}$ by

$$\begin{aligned} u \in P &\Leftrightarrow u_i \sim 0 \text{ for all } i \in I_{\sim} \text{ and } \sim \in \Omega \\ u \in Q &\Leftrightarrow u_i \sim 0 \text{ for all } i \in I^{\sim} \text{ and } \sim \in \Omega. \end{aligned}$$

We lose no generality in assuming that P and Q are of this form. To illustrate, Example 4.4 corresponds to the partitions

$$\begin{array}{lll} I_{<} = \emptyset & I_{=} = \emptyset & I_{\geq} = \{1, 2, 3\} \\ I^{\leq} = \{1\} & I^= = \{3\} & I^{\geq} = \{2\}. \end{array}$$

The point x lies in the common face F , which is the cube in $[-1, 1]^n$ defined by

$$u \in F \quad \Leftrightarrow \quad u_i \begin{cases} \leq 0 & \text{if } i \in I_{<} \cap I^{\leq} \\ \geq 0 & \text{if } i \in I_{\geq} \cap I^{\geq} \\ = 0 & \text{otherwise.} \end{cases}$$

Since $y \in Q$ we deduce

$$y_i \sim 0 \text{ for all } i \in I^{\sim} \text{ and } \sim \in \Omega.$$

It is easy to compute componentwise the *Euclidean* nearest point z' to the point y in the cube P : since $y \in [-1, 1]^n$, for each $i = 1, 2, 3, \dots, n$ we have

$$\begin{aligned} z'_i &= \begin{cases} \min\{y_i, 0\} & (i \in I_{<}) \\ 0 & (i \in I_{=}) \\ \max\{y_i, 0\} & (i \in I_{\geq}) \end{cases} \\ &= \begin{cases} y_i & \text{if } i \in I_{<} \text{ and } y_i \leq 0, \text{ or if } i \in I_{\geq} \text{ and } y_i \geq 0 \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Since $y \in Q$, we deduce

$$z'_i = \begin{cases} y_i & \text{if } i \in \bar{I} \\ 0 & \text{otherwise,} \end{cases}$$

for the index set

$$\bar{I} = (I_{<} \cap I^{\leq}) \cup (I_{\geq} \cap I^{\geq}).$$

Clearly we have $z' \in F \subset Q$, and hence the line segment $[z', y]$ is contained in Q . That line segment is therefore also a geodesic in \mathbf{X} . Since distances between points in \mathbf{X} are no less than their Euclidean distance, we deduce $z' = z$.

Since $[x, w] \subset P$ and $[x, y] \subset Q$, those line segments are also geodesics in \mathbf{X} . The distance in \mathbf{X} between pairs of points, one on each line segment, is never less than the Euclidean distance, so in \mathbf{X} , the angle between those geodesics is never less than their Euclidean angle. Moreover, since the points y, x , and z all lie in the cube Q , and the points z, x , and w all lie in the cube P , the angles $\angle yxz$ and $\angle zxcw$ in \mathbf{X} equal the corresponding Euclidean angles. Thus, it is enough to prove (4.3) for Euclidean angles, and hence, in what follows, we consider all angles to be Euclidean.

Since $x \in F$, the formula above for z shows $\angle yzx = \pi/2$. Consequently we deduce $|x - y|^2 = |x - z|^2 + |z - y|^2$ and

$$\cos(\angle yxz) = \frac{|x - z|}{|x - y|}.$$

Applying the law of cosines, we get

$$\cos(\angle yxw) = \frac{|x - y|^2 + |x - w|^2 - |w - y|^2}{2|x - y||x - w|}$$

and

$$\cos(\angle zxw) = \frac{|x - z|^2 + |x - w|^2 - |w - z|^2}{2|x - z||x - w|}.$$

Thus,

$$\cos(\angle yxw) - \cos(\angle yxz) \cos(\angle zxw) = \frac{|z - y|^2 + |w - z|^2 - |w - y|^2}{2|x - y||x - w|} \leq 0,$$

where the last inequality follows because z is the Euclidean nearest point to y in P and $w \in P$. \square

Consider any distinct points x and a in a CAT(0) cubical complex \mathbf{X} . The CAT(0) property ensures that the distance function to a , denoted by d_a is a convex function on \mathbf{X} [4, Example 2.2.4]. Using the algorithm of [2], we can compute the geodesic $[x, a]$: it passes via a sequence of nontrivial geodesics

$$[x, y_0], [y_0, y_1], [y_1, y_2], \dots, [y_m, a]$$

through some corresponding connected sequence of cells $Q_0, Q_1, Q_2, \dots, Q_{m+1}$ (neither sequence being necessarily unique). We refer to $[x, y_0]$ as an *initial segment* and to Q_0 as a corresponding *initial cell*.

Given a cube $P \subset \mathbf{R}^n$, a vector $v \in \mathbf{R}^n$ is a *subgradient* of a convex function $e: P \rightarrow \mathbf{R}$ at a point $x \in P$ if

$$\langle v, w - x \rangle \leq e(w) - e(x) \quad \text{for all } w \in P.$$

The *subdifferential* $\partial e(x)$ is just the set of such subgradients.

We can now derive our central tool, which allows us to compute subgradients of distance functions restricted to cells in cubical complexes.

Theorem 4.5 (Subgradients of distance functions) *Consider a cube $P \subset \mathbf{R}^n$ and a point $x \in P$. Suppose that P is a cell in a CAT(0) cubical complex \mathbf{X} , and consider any point $a \in \mathbf{X}$. Let d_a^P denote the restriction of the distance function d_a to P . If $a = x$ then $0 \in \partial d_a^P(x)$. Suppose, on the other hand, that $a \neq x$. In*

the geodesic $[x, a]$, let $[x, y]$ be an initial segment corresponding to an initial cell Q . Denote by F the common face of P and Q . In any ambient Euclidean space for Q , denote the nearest point to y in F by z . If $z = x$, then $0 \in \partial d_a^P(x)$. On the other hand, if $z \neq x$, then

$$\frac{\cos(\angle yxz)}{|x-z|}(x-z) \in \partial d_a^P(x).$$

Proof The case $a = x$ is trivial, so assume henceforth $a \neq x$. By definition, we have $y \neq x$.

Consider first the case $z = x$. In that case, by Theorem 4.2, x is the nearest point to y in P , and since $y \in [x, a]$, we deduce that x is also the nearest point to a in P (see [9, Proposition II.2.4]), so $0 \in \partial d_a^P(x)$ follows.

On the other hand, consider the case $z \neq x$. We must prove

$$\frac{\cos(\angle yxz)}{|x-z|} \langle x-z, w-x \rangle \leq d_a^P(w) - d_a^P(x) \quad \text{for all } w \in P.$$

This inequality holds when $w = x$. When $w \neq x$, we have

$$\frac{d(a, x) - d(a, w)}{d(x, w)} \leq \cos(\angle axw) = \cos(\angle yxw) \leq \cos(\angle yxz) \cdot \cos(\angle zxw),$$

by Lemma 4.1 and Theorem 4.2. The result now follows. \square

To illustrate the result, we consider a simple example.

Example 4.6 (Calculating a subgradient of a distance function)

Consider the CAT(0) cubical complex \mathbf{X} the cells of which are the following squares in \mathbf{R}^2 , along with their edges and vertices:

$$\begin{aligned} P &= [0, 1] \times [0, 1] & Q &= [0, 1] \times [-1, 0] \\ S &= [0, 1] \times [-2, -1] & T &= [-1, 0] \times [-2, -1]. \end{aligned}$$

For the point $a = (-\frac{1}{2}, -2) \in \mathbf{X}$, the distance function $d_a: \mathbf{X} \rightarrow \mathbf{R}$ is given by

$$d_a(w) = \begin{cases} \frac{1}{2}\sqrt{5} + \sqrt{w_1^2 + (w_2 + 1)^2} & (w_1 \geq 0 \text{ and } w_2 \geq 2w_1 - 1) \\ \sqrt{(w_1 + \frac{1}{2})^2 + (w_2 + 2)^2} & (\text{otherwise}). \end{cases}$$

This function must be convex on \mathbf{X} , so in particular it is convex on $P \cup Q$. Differentiating shows that, for points w in the interior of the set \mathbf{X} , regarded as a subset of \mathbf{R}^2 , we have

$$\nabla d_a(w) = \begin{cases} \frac{1}{\sqrt{w_1^2 + (w_2 + 1)^2}}(w_1, w_2 + 1) & (w_1 > 0 \text{ and } w_2 > 2w_1 - 1) \\ \frac{1}{\sqrt{(w_1 + \frac{1}{2})^2 + (w_2 + 2)^2}}(w_1 + \frac{1}{2}, w_2 + 2) & (w_1 < 0 \text{ or } w_2 < 2w_1 - 1). \end{cases}$$

The two cases in this formula describe a set of full measure around the point $x = (\frac{1}{2}, 0)$, and as w approaches x within this set, the gradient has a unique limit: $\frac{1}{\sqrt{5}}(1, 2)$. Standard properties of Euclidean convex functions therefore show that the function d_a restricted to $P \cup Q$ has Gateaux derivative $\frac{1}{\sqrt{5}}(1, 2)$ at x . The Euclidean normal cone to P at x is $\{0\} \times -\mathbf{R}_+$, so we deduce

$$(4.7) \quad \partial d_a^P(x) = \frac{1}{\sqrt{5}}(1, 2) + (\{0\} \times -\mathbf{R}_+).$$

We now compare this with the subgradient provided by Theorem 4.5. Let y denote the point $(0, -1)$. The geodesic from x to a consists of the two line segments $[x, y]$ in the square Q and $[y, a]$ in the square T . The common face of P and Q is $F = [0, 1] \times \{0\}$, and in the square Q , the nearest point to y in the face F is the point $z = (0, 0)$. The unit vector in the direction $x - z$ is therefore $(1, 0)$. The angle $\text{angle}xyz$ has cosine $\frac{1}{\sqrt{5}}$, so the theorem asserts that the vector $\frac{1}{\sqrt{5}}(1, 0)$ is a subgradient of d_a^P at x . This is confirmed by equation (4.7).

5 Means, medians, and circumcenters

In any metric space (\mathbf{X}, d) , we can consider a variety of interesting optimization problems associated with a given nonempty finite set $A \subset \mathbf{X}$, posed simply in terms of the distance functions d_a to points $a \in A$. Foremost among these are *weighted mean* problems, which entail minimizing functions $f: \mathbf{X} \rightarrow \mathbf{R}$ defined by

$$(5.1) \quad f(x) = \sum_{a \in A} \lambda_a d_a^q(x) \quad (x \in \mathbf{X})$$

for some given weights $\lambda_a > 0$, for $a \in A$, and a given exponent $q \geq 1$. The special case $q = 1$ is known as the *median* problem. Another well-known example is the *minimum enclosing ball* or *circumcenter* problem, which entails minimizing over \mathbf{X} the function

$$x \mapsto \max_{a \in A} d_a(x).$$

When the space \mathbf{X} is a finite CAT(0) cubical complex, we can solve weighted mean and minimum enclosing ball problems using Algorithm 3.2 and Theorem 4.5. Consider first the weighted mean problem [4, Example 6.3.4]. The function (5.1) is convex, and a minimizer always exists, by [4, Lemma 2.2.19]. For $q > 1$, the function f is strictly convex, so the minimizer is unique. In the case $q = 1$, minimizers may not be unique.

To compute means and medians in general Hadamard spaces, the only methods previously analyzed appear to be the splitting proximal point algorithms described in [6]. As we noted in the introduction, such methods are inevitably slow. In the

case of a finite CAT(0) cubical complex, we instead propose Algorithm 3.2 as a faster alternative.

To implement Algorithm 3.2, we first fix attention on some cell P in the complex \mathbf{X} . Via an isometric embedding, we make the identification $P = [0, 1]^n \subset \mathbf{R}^n$. Using the notation of Theorem 4.5, we seek to minimize the convex function

$$f_P = \sum_{a \in A} \lambda_a (d_a^P)^q: P \rightarrow \mathbf{R}.$$

For this purpose, we can use a standard Euclidean cutting plane method, of the kind described in [11, Chapter 2]. Such methods require a separation/subgradient oracle of the following form. Consider any input point $x \in \mathbf{R}^n$.

- If $x \notin P$, then output a hyperplane separating x from P .
- Otherwise, return the value $f_P(x)$ and a subgradient $v \in \partial f_P(x)$.

Theorem 4.5 allows us to implement this oracle. Separating any point from P is trivial. On the other hand, for any point x in P , for each point $a \in A$, the theorem describes how to calculate a corresponding subgradient $v_a \in \partial d_a^P(x)$. We deduce $q d_a^{q-1}(x) v_a \in \partial ((d_a^P)^q)(x)$, and then adding gives a subgradient:

$$\sum_a \lambda_a q d_a^{q-1}(x) v_a \in \partial f_P(x).$$

Example 5.2 (A simple mean calculation, concluded) We return to the simple mean problem in Examples 2.1 and 3.1. On each of the three cells we solve the corresponding subproblem using a standard subgradient-based method, deploying the subgradients available through Theorem 4.5. We compare six methods: four cutting plane methods, the classical subgradient method, and the cyclic proximal point method from [6]. The cutting plane methods consist of the classical ellipsoid algorithm and three standard bundle methods from [20]: a level bundle method, a proximal bundle method, and a dual level method. The methods all converge to the optimal solution of each subproblem: the true mean in the cell P_1 , and the point $(0, 0)$ in the other two cells. The behavior of the objective value in the case of the cell P_1 is shown in Figure 1, plotted against the number of geodesics computed: the behavior in the other cells is similar. Even on this small low-dimensional example, the cyclic proximal point method is relatively slow. The subgradient method, although clearly sublinear, converges reasonably quickly on this example, in part because the objective function happens to be smooth around the optimal solution. The cutting plane methods are all faster, and the plot suggests the linear convergence we expect.

Turning to the minimum enclosing ball problem [4, Example 2.2.18], we can frame the problem as minimizing the strongly convex function $f: \mathbf{X} \rightarrow \mathbf{R}$ defined by

$$f(x) = \max_{a \in A} d_a^2(x) \quad (x \in \mathbf{X}).$$

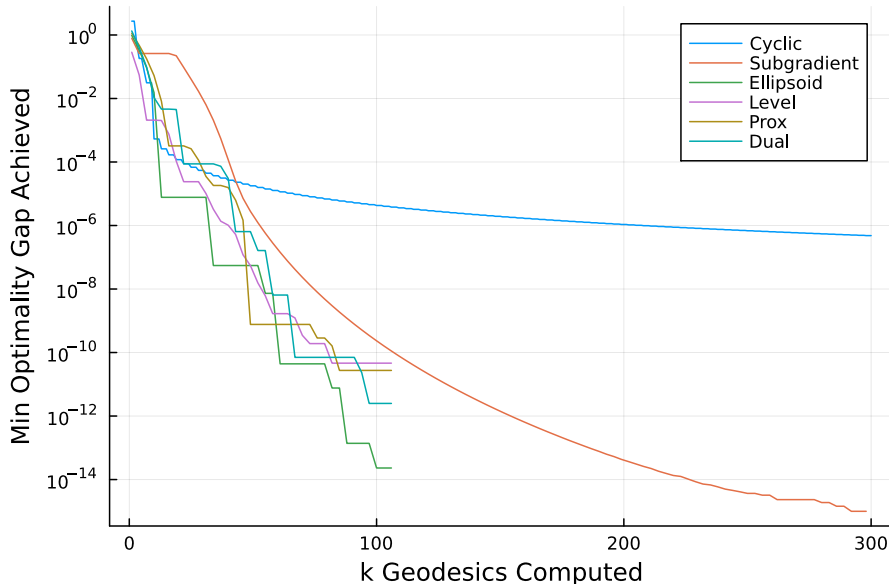


Figure 1: A simple mean calculation. Five subgradient-based methods for minimizing, over a single cell, the sum of the squared distances to three points in a complex of three squares, compared with the cyclic proximal point method. Each graph plots the smallest objective value found so far against the number of geodesics computed.

The unique solution is called the *circumcenter* of A . Previous literature makes no apparent reference to algorithms for the minimum enclosing ball problem in general Hadamard spaces. A subgradient-style method has been proposed recently in [21], with the slow convergence rate typical of subgradient methods. In the case of CAT(0) cubical complexes, we again propose Algorithm 3.2 as a faster alternative.

We follow the same approach, focusing on some cell P and minimizing the strongly convex function $f_P: P \rightarrow \mathbf{R}$ defined by

$$f_P(x) = \max_{a \in A} (d_a^P)^2(x) \quad (x \in \mathbf{X}).$$

To compute a subgradient of f_P at a point $x \in P$, we first select, from the set A , a point a at maximum distance from x . We then appeal to Theorem 4.5 to calculate a corresponding subgradient $v_a \in \partial d_a^P(x)$. Standard Euclidean subgradient calculus [27] now gives us the desired subgradient: $2d_a(x)v_a \in \partial f_P(x)$.

To apply Algorithm 3.2 to a mean, median, or circumcenter problem, having implemented the separation/subgradient oracle, we next choose a cutting plane method to solve the subproblems over cells P . Most classically — albeit not the best choice in theory or practice, as discussed in the introduction — we could apply the ellipsoid algorithm, starting with the smallest Euclidean ball containing P . That algorithm converges linearly: by [11, Theorem 2.4], if the cell P has dimension n , then after

$t \geq n^2 \log n$ calls to the oracle, we find an iterate in P at which the value of f exceeds $\min_P f$ by no more than

$$2\sqrt{n}(\max_P f) \exp\left(-\frac{t}{2n^2}\right).$$

If f is in fact *strongly* convex (as in the case of the mean or circumcenter problems), then the *iterates* also converge linearly, to the unique minimizer of f on P .

6 The distance to an intersection of balls

In addition to mean and circumcenter problems, another fundamental problem associated with a given nonempty finite subset A of a metric space (\mathbf{X}, d) , and posed simply in terms of distance functions, is the *intersecting balls* problem. Centered at each point $a \in A$, given a radius $\rho_a > 0$, we consider the closed ball $B_{\rho_a}(a)$. We seek a point in their intersection:

$$\mathcal{F} = \bigcap_{a \in A} B_{\rho_a}(a).$$

In any Hadamard space \mathbf{X} , we could approach the intersecting balls problem via the method of cyclic projections [3, 5]: the projection of any point x onto a ball centered at a is easy to compute from the geodesic $[x, a]$, since the projection must lie on the geodesic. However, in a general Hadamard space, the rate at which the method converges to a point in the intersection is not clear, and when the intersection is empty, the behavior of the iterates is not well understood [23].

When the space \mathbf{X} is a finite CAT(0) cubical complex, we take a different approach to the intersecting balls problem, again using Algorithm 3.2 and Theorem 4.5. Consider the related problem of minimizing over \mathbf{X} the objective function

$$\sum_{a \in A} \text{dist}_{B_{\rho_a}(a)}.$$

Since this function is coercive, it has a minimizer, by [4, Lemma 2.2.19]. Consider any such minimizer \bar{x} . If $\bar{x} \in \mathcal{F}$, then we have solved our problem, and on the other hand, if $\bar{x} \notin \mathcal{F}$, then the set \mathcal{F} must be empty. We therefore seek to minimize the function $f: \mathbf{X} \rightarrow \mathbf{R}$ defined by

$$f(x) = \sum_{a \in A} \max\{d_a(x) - \rho_a, 0\} \quad (x \in \mathbf{X}).$$

We follow the same strategy as for the mean problem in the previous section, focusing on individual cells P in the complex \mathbf{X} , making the identification $P = [0, 1]^n$ and seeking to minimize the convex function $f_P: P \rightarrow \mathbf{R}$ defined by

$$f_P(x) = \sum_{a \in A} \max\{d_a^P(x) - \rho_a, 0\} \quad (x \in P).$$

As before, to implement a cutting plane algorithm, we need to find a subgradient of f_P at any input point $x \in P$. To that end, for each point $a \in A$, we compute the distances $d_a(x)$, and calculate a vector v_a as follows. If $d_a(x) \leq \rho_a$, then we define $v_a = 0$. On the other hand, if $d_a(x) > \rho_a$, then we appeal to Theorem 4.5 to calculate a subgradient $v_a \in \partial d_a^P(x)$. In that way, standard Euclidean subgradient calculus [27] guarantees that v_a is a subgradient of the convex function given by

$$x \mapsto \max\{d_a^P(x) - \rho_a, 0\} \quad (x \in P),$$

from which we deduce our desired subgradient:

$$\sum_{a \in A} v_a \in \partial f_P(x).$$

We can now apply a cutting plane algorithm.

Refining the intersecting balls problem, we might seek to compute the distance from a given point $b \in \mathbf{X}$ to the intersection of balls \mathcal{F} . We can approach this problem by a bisection search strategy. We first run the previous algorithm, either finding a point $\hat{x} \in \mathcal{F}$ or detecting that none exists, in which case we terminate. If we find $\hat{x} \in \mathcal{F}$, we begin our bisection search, always maintaining an interval $[\underline{\rho}, \bar{\rho}]$ containing the distance from b to \mathcal{F} . Initially we set $\underline{\rho} = 0$ and $\bar{\rho} = d(b, \hat{x})$. At each iteration, we consider the midpoint ρ of the current interval, and apply the previous algorithm to the intersecting balls problem $\mathcal{F} \cap B_\rho(b)$. If the intersection is empty, we update $\underline{\rho} = \rho$; otherwise we update $\bar{\rho} = \rho$. We then repeat.

References

- [1] A. Abrams and R. Ghrist. State complexes for metamorphic robots. *The International Journal of Robotics Research*, 23:811–826, 2004.
- [2] F. Ardila, M. Owen, and S. Sullivant. Geodesics in CAT(0) cubical complexes. *Advances in Applied Mathematics*, 48:142–163, 2012.
- [3] D. Ariza-Ruiz, G. López-Acedo, and A. Nicolae. The asymptotic behavior of the composition of firmly nonexpansive mappings. *J. Optim. Theory Appl.*, 167:409–429, 2015.
- [4] M. Bačák. *Convex Analysis and Optimization in Hadamard Spaces*. De Gruyter, Berlin, 2014.
- [5] M. Bačák, I. Searston, and B. Sims. Alternating projections in CAT(0) spaces. *J. Math. Anal. Appl.*, 385:599–607, 2012.
- [6] M. Bačák. Computing medians and means in Hadamard spaces. *SIAM Journal on Optimization*, 24:1542–1566, 2014.

- [7] L.J. Billera, S.P. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Advances in Applied Mathematics*, 27:733–767, 2001.
- [8] N. Boumal. *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, Cambridge, 2023.
- [9] M.R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer-Verlag, Berlin, 1999.
- [10] D.G. Brown and M. Owen. Mean and variance of phylogenetic trees. *Systematic Biology*, 69:139–154, 2020.
- [11] S. Bubeck. Convex optimization: algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- [12] R. Ghrist and V. Peterson. The geometry and topology of reconfiguration. *Advances in Applied Mathematics*, 38:302–323, 2007.
- [13] A.J. Goldman. Minimax location of a facility in a network. *Transportation Science*, 6:407–418, 1972.
- [14] M. Gromov. *Structures métriques pour les variétés Riemanniennes*, volume 1 of *Textes Mathématiques*. CEDIC, Paris, Edited by J. Lafontaine and P. Pansu, 1981.
- [15] P. Hansen, M. Labbé, D. Peeters, and J.-F. Thisse. Single facility location on networks. *Annals of Discrete Mathematics*, 31(113-146), 1987.
- [16] K. Hayashi. A polynomial time algorithm to compute geodesics in CAT(0) cubical complexes. *Discrete and Computational Geometry*, 65:636–654, 2021.
- [17] J. Holbrook. No dice: a deterministic approach to the Cartan centroid. *Journal of the Ramanujan Mathematical Society*, 27:509–521, 2012.
- [18] T. Holz, S. Huckemann, Huiling Le, J.S. Marron, J.C. Mattingly, E. Miller, J. Nolen, M. Owen, V. Patrangenaru, and S. Skwerer. Sticky central limit theorems on open books. *Annals of Applied Probability*, 23:2238–2258, 2013.
- [19] Yin Tat Lee, A. Sidford, and S.S. Vempala. Efficient convex optimization with membership oracles. *Proceedings of Machine Learning Research*, 75:1–3, 2018 [arXiv:1706.07357](#).
- [20] C. Lemaréchal, A. Nemirovskii, and Y. Nesterov. New variants of bundle methods. *Math. Programming*, 69(1, Ser. B):111–147, 1995. Nondifferentiable and large-scale optimization (Geneva, 1992).
- [21] A.S. Lewis, G. López-Acedo, and A. Nicolae. Horoballs and the subgradient method. [arXiv:2403.15749](#), 2024.
- [22] Yongdo Lim and M. Palfia. Weighted deterministic walks for the least squares mean on Hadamard spaces. *Bulletin of the London Mathematical Society*, 46:561–570, 2014.

- [23] A. Lytchak and A. Petrunin. Cyclic projections in Hadamard spaces. *J. Optim. Theory Appl.*, 194:636–642, 2022.
- [24] E.M. Massart, J.M. Hendrickx, and P.-A. Absil. Matrix geometric means based on shuffled inductive sequences. *Linear Algebra and its Applications*, 542:334–359, 2018.
- [25] E. Miller, M. Owen, and J.S. Provan. Polyhedral computational geometry for averaging metric phylogenetic trees. *Advances in Applied Mathematics*, 68:51–91, 2015.
- [26] M. Owen and S. Provan. A fast algorithm for computing geodesic distances in tree space. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 8:2–13, 2011.
- [27] R.T. Rockafellar. *Convex analysis*. Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J., 1970.
- [28] K.-T. Sturm. Nonlinear martingale theory for processes with values in metric spaces of nonpositive curvature. *Ann. Probab.*, 30:1195–1222, 2002.
- [29] B.C. Tansel, R.L. Francis, and T.J. Lowe. State of the art — location on networks: a survey. Part I: the p-center and p-median problems. *Management Science*, 29:482–497, 1983.
- [30] P.M. Vaidya. A new algorithm for minimizing convex functions over convex sets. *Mathematical Programming*, 73:291–341, 1996.
- [31] Hongyi Zhang and S. Sra. First-order methods for geodesically convex optimization. *Journal of Machine Learning Research*, 49:1–22, 2016.

A Appendix: core cells

An interesting source of illustrative computational examples arises from the class of cubical complexes \mathbf{X} consisting of a collection of cells all sharing one common face. As observed in [18], although this face is lower-dimensional than \mathbf{X} it often contains the mean of a set of points distributed throughout \mathbf{X} .

Example A.1 (An open book [18]) Consider three copies of the square $[0, 1]^2$ joined along the shared edge $\{0\} \times [0, 1]$: the “spine” of the open book. Denote the three copies P^1, P^2, P^3 , and consider any points $a^i \in \text{int } P^i$ satisfying

$$a^i < \sum_{j \neq i} a^j \quad \text{for } i = 1, 2, 3.$$

(An example is $a^i = (\frac{1}{2}, \frac{1}{2})$ for each i .) A quick calculation shows that the mean is the following point on the spine:

$$\left(0, \frac{1}{3}(a_2^1 + a_2^2 + a_2^3)\right).$$

The spine is thus sticky: the mean remains there for all small perturbations of the three points.

Definition A.2 In a cubical complex, if all distinct pairs of maximal cells have the same nonempty intersection F , then we call F the *core*.

The core, if it exists, is a common face of all the maximal cells. For example, the 3-spider has the core $\{0\}$. The maximal cells may not all have the same dimension. For example, the complex consisting of the maximal cells in \mathbf{R}^3

$$[0, 1] \times [0, 1] \times [0, 1] \quad \text{and} \quad [-1, 0] \times \{0\} \times [0, 1] \quad \text{and} \quad \{0\} \times [-1, 0] \times [0, 1]$$

has core $\{0\} \times \{0\} \times [0, 1]$.

Proposition A.3 *Any cubical complex with a core is CAT(0).*

Proof Consider a cubical complex \mathbf{X} with a core F . The cells of \mathbf{X} consist just of all the faces of maximal cells. Since \mathbf{X} is contractible to any point in F , it is simply connected.

We next observe that Gromov's link condition holds. To see this, given any vertex, consider any finite set E of k incident edges, each pair of which lie in a common 2-dimensional cell of \mathbf{X} . For each maximal cell P , let E_P denote those edges in E that lie in P but not in the common face F . Consider any two maximal cells P, Q . Suppose that there exist edges $p \in E_P \setminus E_Q$ and $q \in E_Q \setminus E_P$. Any 2-dimensional cell containing both p and q must be a face of some maximal cell R . The edge p lies in both R and P , but by assumption not in F , so in fact $R = P$. But the same argument shows $R = Q$, giving a contradiction. We deduce that either $E_P \subset E_Q$ or $E_Q \subset E_P$, proving that the collection of sets E_P for maximal cells P is totally ordered by inclusion. The collection therefore has maximal element, E_P for some maximal cell P . Every edge in E lies in this cell P , and hence in some k -dimensional face of P , as required. \square

We next compute geodesics.

Proposition A.4 *Consider points x and y in a cubical complex \mathbf{X} with a core F . If some cell contains both points, then the geodesic $[x, y]$ is the line segment joining them in that cell. On the other hand, if no cell contains both, then there are maximal cells P and Q such that $x \in P \setminus Q$ and $y \in Q \setminus P$. In P , consider the distance $|x - x_F|$ between x and x_F , its orthogonal projection onto F . Analogously, in Q , consider the distance $|y - y_F|$ between y and y_F , its orthogonal projection onto F . Consider the following convex combination in F :*

$$z = \frac{1}{|x - x_F| + |y - y_F|} (|y - y_F| x_F + |x - x_F| y_F).$$

Then the geodesic $[x, y]$ is the union of the two geodesics $[x, z]$ and $[z, y]$.

Proof Note that the geodesic $[x, y]$ must pass through the face F , so it is the union of two geodesics $[x, w]$ and $[w, y]$ for some point $w \in F$. The length of this path is

$$\sqrt{|w - x_F|^2 + |x - x_F|^2} + \sqrt{|w - y_F|^2 + |y - y_F|^2}.$$

As w varies over F , this function is convex, and viewed on any ambient Euclidean space in which F lives, its gradient is

$$(1.5) \quad \frac{1}{\sqrt{|w - x_F|^2 + |x - x_F|^2}}(w - x_F) + \frac{1}{\sqrt{|w - y_F|^2 + |y - y_F|^2}}(w - y_F).$$

Notice

$$z - x_F = \frac{|x - x_F|}{|x - x_F| + |y - y_F|}(y_F - x_F),$$

so

$$\sqrt{|z - x_F|^2 + |x - x_F|^2} = \frac{|x - x_F|}{|x - x_F| + |y - y_F|} \sqrt{|y_F - x_F|^2 + (|x - x_F| + |y - y_F|)^2}$$

and hence the first term in the expression (1.5), when $w = z$, becomes

$$\frac{1}{\sqrt{|y_F - x_F|^2 + (|x - x_F| + |y - y_F|)^2}}(y_F - x_F).$$

By symmetry, the second term is identical, after multiplying by -1 . Hence the gradient (1.5) is zero when $w = z$, so z indeed minimizes the length. \square