

Bumpy Pyramid Folding

Zachary R. Abel*

Erik D. Demaine†

Martin L. Demaine‡

Hiro Ito§

Jack Snoeyink¶

Ryuhei Uehara||

Abstract

We investigate folding problems for a class of petal polygons P , which have an n -polygonal base B surrounded by a sequence of triangles. We give linear time algorithms using constant precision to determine if P can fold to a pyramid with flat base B , and to determine a triangulation of B (crease pattern) that allows folding into a convex (triangulated) polyhedron. By Alexandrov’s theorem, the crease pattern is unique if it exists, but the general algorithm known for this theorem is pseudo-polynomial, with very large running time; ours is the first efficient algorithm for Alexandrov’s theorem for a special class of polyhedra. We also give a polynomial time algorithm that finds the crease pattern to produce the maximum volume triangulated polyhedron.

1 Introduction

In 1525, the German painter Albrecht Dürer published his masterwork on geometry [7], whose title translates as, “On Teaching Measurement with a Compass and Straightedge for lines, planes, and whole bodies.” In the book, he presented each polyhedron by drawing a *net* for it: an unfolding of the surface to a planar layout. To this day it remains a big open problem whether every convex polyhedron has a (non-overlapping) net. Several strategies and algorithms for unfolding a polyhedron have been investigated. One, called *star unfolding*, is produced by cutting along every shortest path to each vertex¹.

To understand unfolding, it is interesting to look at the inverse: one folding problem asks what polyhedra can be folded from a given polygonal sheet of paper. For




Figure 1: A petal polygon and resulting pyramid.

example, the Latin cross, which is a typical unfolding of a cube, can form 23 polyhedra by 85 distinct ways of folding. Comprehensive surveys of folding and unfolding can be found in [6, 9].

In this paper, we investigate a folding problem for a class of polygons that can come from star unfolding, or from unfolding pyramids by cutting all edges incident on an apex. A polygon $P = (p_1, c_1, p_2, c_2, \dots, p_n, c_n)$ is a *petal polygon* if and only if it satisfies two conditions.

1. The subpolygon $B = (p_1, p_2, \dots, p_n)$, called the *base* of the polygon P , is convex and each point c_i is outside of B .
2. The pairs of edges incident on each p_i have equal length. (That is, for each $i = 1, 2, \dots, n$, lengths $|p_i c_i| = |c_{i-1} p_i| = \ell_i^2$.)

We investigate the *petal folding problem*: determine if some specific polyhedra can be folded from a petal polygon P by gluing the pairs of edges incident on each p_i (all vertices c_i meet at the *apex* c). We investigate three specific variations of the petal folding problem.

First we consider the conditions under which the n -gon base B can be kept flat. This *petal pyramid folding* problem asks whether we can obtain the pyramid from a given petal polygon. This problem can be solved in linear time.

Theorem 1 *Let P be a petal polygon. Then the petal pyramid folding problem can be solved in linear time.*

²In this paper, all indices are computed mod n .

*Department of Mathematics, MIT, zabel@math.mit.edu
 †Computer Science and Artificial Intelligence Laboratory, MIT, edemaine@mit.edu
 ‡Computer Science and Artificial Intelligence Laboratory, MIT, mdemaine@mit.edu
 §School of Informatics and Engineering, The University of Electro-Communications, Japan, itohiro@uec.ac.jp
 ¶Department of Computer Science, The University of North Carolina, USA, snoeyink@cs.unc.edu
 ||School of Information Science, JAIST, Japan, uehara@jaist.ac.jp

¹In this paper, we only concern with starting and final configurations. In other word, we do not insist that if there is crossing throughout the folding motion or not.

When we cannot fold a petal polygon P into a pyramid, we may still be able to fold P into a polyhedron, collecting the vertices c_i to an apex c , by triangulating the base B to make a pattern of creases. A *crease pattern* is a planar straight-line graph with an assignment of mountain or valley to each edge. From the perspective of an observer, a *valley fold* has dihedral angle $< 180^\circ$, and a *mountain fold* has dihedral angle $> 180^\circ$.

A *bumpy pyramid* is a polyhedron obtained by folding of a petal polygon P that collects the vertices c_i to an apex c by folding the base B according to some crease pattern on a triangulation of B (if it exists). We will show later that once we fix a triangulation of the base B , its feasible crease pattern (i.e., mountain/valley assignments on the triangulation) to fold a bumpy pyramid is uniquely determined if it exists. However, finding some specific triangulation with its feasible crease pattern can be surprisingly complex. Even for $n = 4$, a petal polygon P typically folds into two different bumpy pyramids (one convex and one non-convex), but sometimes can fold to only one or even to no bumpy pyramid.

Thus, for the second problem, we add the condition of convexity and define *convex bumpy pyramid folding*: find a crease pattern to fold a convex bumpy pyramid, or determine that no such crease pattern exists. Here the crease pattern is a triangulation of the base B with all creases assigned to be mountain folds. Through a connection to power diagrams [3], every P can fold into a unique convex polyhedron if the lengths ℓ_i are sufficiently long, and the corresponding crease pattern (triangulation) on B can be found in linear time using constant precision:

Theorem 2 *Let P be a petal polygon with $2n$ vertices. Then the convex bumpy pyramid folding problem can be solved in linear time.*

Two things suggested by the wording of this theorem help us. First, once the triangulation is chosen, there is little further choice for the crease pattern — all creases must be flat or mountain folds for the result to be convex. Second, Alexandrov’s theorem implies that, if a petal polygon P can fold into a convex polyhedron, then that folding is unique. Thus, the problem is to find the unique crease pattern on B — in other words, the triangulation of B — so that P folds into a convex polyhedron by mountain folding on the creases.

Let us elaborate on this, and its related results. 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. Thus, if P is a net of a convex bumpy pyramid, then the shape (as a convex polyhedron) is uniquely determined. Alexandrov’s theorem was stated in 1942, and a constructive proof was given by Bobenko and Izestiev in 2008 [4]. A pseudo-polynomial algorithm for Alexandrov’s theorem, given by Kane *et al.* in

2009, runs in $O(n^{456.5r^{1891}}/\epsilon^{121})$ time, where r is the ratio of the largest and smallest distances between vertices, and ϵ is the coordinate relative accuracy [8]. The exponents in the time bound of the result are remarkably huge. As far as the authors know, Theorems 1 and 2 are the first efficient algorithms for Alexandrov’s theorem for a family of nontrivial convex polyhedra.

Aronov and O’Rourke used a Voronoi diagram to prove that the star unfolding has no overlap in the plane [2]. In star unfolding, a base face remains flat and flaps are obtained by straightening cuts along geodesic paths that cross several faces to reach the same point c . The sites for this Voronoi diagram are the copies of c , which would be the c_i s, and the metric is Euclidean distance in the plane. In convex bumpy pyramid folding, as we will see, the triangular flaps remain flat, while the base is folded along diagonals. These diagonals will form a triangulation dual to the power diagram, which is a Voronoi diagram of sites p_i using the power distance.

Finally, we turn to the third problem, which is to find a crease pattern of B such that P folds into the bumpy pyramid of *maximum* volume. At first glance, it may seem that the convex bumpy pyramid should achieve the maximum volume, but this is not true in general. In this paper, we give counterexamples: two petal polygons whose non-convex bumpy pyramids have larger volumes than their convex ones. We here note that, if the crease pattern of B is fixed, then its volume can be computed by decomposing the polyhedron into tetrahedra by cuts from the apex to the creases and summing the volumes of the tetrahedra. Even though the number of possible triangulations of the base B can be exponential, we can find the bumpy pyramid of maximum volume by nontrivial dynamic programming in $O(n^3)$ time:

Theorem 3 *Let P be a petal polygon whose base B has n vertices. The crease pattern (triangulation of B) that gives the bumpy pyramid with maximum volume can be found in $O(n^3)$ time.*

2 Preliminaries

We sometimes discuss the convexity of a polyhedron folded from a petal polygon P . We may have a special case that the polyhedron is concave at its apex c . We first consider this special case. When the total sum of interior angles at the c_i is greater than 360° and $n > 3$, the apex will be a point of negative curvature, and must be a saddle in any folded polyhedron. A simple example is drawn in Figure 2: by mountain folding around the edges of B and p_1p_3 , we have a polyhedron that is a saddle at the apex c . The case $n = 3$ is exceptional; when the total sum of interior angles at the c_i is greater than 360° , ℓ_i s are too short to make a polyhedron.

In our first two problems, the solutions require that the result be a convex folded polyhedron, which is pos-




Figure 2: This net forces a saddle point at apex c .




Figure 3: The locus of $c'_i(t)$ for folding $\triangle p_i p_{i+1} c_i(t)$.

sible only if the curvature at the apex is positive. Since it can be checked in linear time, we basically eliminate the case to make our arguments clear. We need a little more notation to describe folding processes. We suppose initially that P is on the xy plane, and that each point p of P is given by Cartesian coordinates in the plane, namely $(x(p), y(p))$. When we consider the folding of P in general, we put P into an xyz space, and then each point p of P is described by three coordinates, $(x(p), y(p), z(p))$. For two points p and q in the plane or in space, the line that contains the line segment pq is denoted by \overline{pq} .

Let T_i denote the triangle formed by the three points $p_i p_{i+1} c_i$. Consider a folding of the triangle T_i onto B along the line $\overline{p_i p_{i+1}}$. In the plane this would be a reflection or *flip* across the line; in space the movement is described as the function of time, rotating the triangle about $\overline{p_i p_{i+1}}$ from the initial position at $t = 0$ back into the plane at $t = 1$ (Figure 3). That is, when we flip the triangle T_i , the point c_i moves from $c_i(0)$ to $c_i(1)$; if c_i stops at apex c and is glued to form a (bumpy) pyramid, we have $c = c_i(t)$ for some $0 < t < 1$.

We can also watch the projection $c'_i(t)$ of the motion of c_i back in the xy plane. That is, for $c_i = (x(c_i), y(c_i), z(c_i))$, the projection $c'_i = (x(c_i), y(c_i), 0)$. We define τ_i , the *trace* of c_i , as the locus of all projections c'_i as c_i moves from $c_i(0)$ to $c_i(1)$ through space. In fact, τ_i is simply the segment $c'_i(0)c'_i(1)$, which lies in the xy plane; we will use τ_i when we are focusing on the plane, and $c_i(0)c_i(1)$ when we are thinking of this as the limits of motion of c_i in space.

Observation 1 *We suppose that we mountain fold the triangle T_i onto B by flipping it. The trace τ_i of $c_i(t)$ is the chord between the intersections of the two circles,*

C_i and C_{i+1} , centered at p_i and p_{i+1} and with radius $|p_i c_i|$ and $|p_{i+1} c_i|$, respectively. The line through τ_i is perpendicular to the line $\overline{p_i p_{i+1}}$.

Note that the trace is easily constructed by reflection about a line in the plane.

3 Folding to Pyramids

In this section, we suppose that the base B is on an xy plane, and each point c_i moves to the positive side of B , and generates a trace τ_i on the xy plane. (We observe from outside, so this folding is *mountain folding* along the line $p_i p_{i+1}$ for some i . That is, we always have $p_i = (x(p_i), y(p_i), 0)$ and $c_i = (x(c_i), y(c_i), z(c_i))$ with $z(c_i) \geq 0$. The opposite folding is *valley folding*.)

The c_i can meet at a common apex c if and only if the corresponding traces, τ_i , meet at a common point in the plane. In fact, it is enough if $n - 1$ of the traces meet at a common point, because then the n th will do as well.

Lemma 4 *Let P be a petal polygon whose base B has n sides. Then P can fold to a pyramid if and only if the $n - 1$ traces τ_i intersect at the same point c' .*

Proof. Suppose that the points p_i lie in the xy plane in three dimensions. Consider spheres S_i centered at p_i of radius ℓ_i , which makes S_i pass through c_{i-1} and c_i . Each trace τ_i is the projection onto the xy plane of the disk bounded by the intersection $S_i \cap S_{i+1}$. If $n - 1$ traces intersect in a common point c' , then all spheres contain c' , so it is in the remaining trace as well.

Now, if P folds to a pyramid, then all c_i s must meet at a common apex c . That is, for all $i \in [1..n]$, there exists an $\alpha_i \in (0, 1)$ with $c_i(\alpha_i) = c$. By Observation 1, the projection c' will lie on all traces τ_i .

For the reverse, if a point c' lies on trace τ_i , then there is a unique $\alpha_i \in (0, 1)$ with projection $c'_i(\alpha_i) = c'$. Thus, $c_i(\alpha_i)$ on the intersection $S_i \cap S_{i+1}$, which have the same height above c' . Similarly, $c_{i-1}(\alpha_{i-1})$ lies on the intersection $S_{i-1} \cap S_i$, so these two spheres have the same height above c' . By transitivity, all spheres have the same height above c' , so all $c_i(\alpha_i)$ meet at the same point c , which is the apex c of a folded pyramid. \square

Now, to prove Theorem 1, it is sufficient to show that we can check for a common intersection c' of traces in linear time. The candidate position is the solution to simultaneous linear equations $(c' - c_i) \cdot (p_i - p_{i+1}) = 0$. By Lemma 4, it is enough to look at $n - 1$ equations, so for $n = 3$ always has a unique candidate; for $n > 3$ we check that the solution for 2 independent equations satisfies the rest. Once we have the candidate c' , we simply check that its distance to the line $\overline{p_i p_{i+1}}$ is less than that of c_i . Using $(x, y)^\perp = (-y, x)$, this is

$$|(p_i - p_{i+1})^\perp \cdot (c' - p_i)| \leq (p_i - p_{i+1})^\perp \cdot (c_i - p_i).$$

Using Cramer's rule to find c' shows that the inequalities can be evaluated with four times the input precision, so constant time apiece, and $O(n)$ time overall. This completes the proof of Theorem 1.

We note that for any base B and chosen origin o , both in the xy plane, by making the triangle flaps to the c_i sufficiently long, it is always possible to fold a pyramid that has its apex above o . (Actually, you only get to choose the length of one flap, because then all the other edge lengths, as well as the height of the apex are determined.) In proving this we introduce the power distance.

Theorem 5 *Let $B = (p_1, p_2, \dots, p_n)$ be a convex polygon in a Cartesian plane with origin o . If we choose any one length $l_i > |p_i o|$, it is possible to construct the points (c_1, c_2, \dots, c_n) for a petal polygon that folds to a pyramid with apex c projecting onto the origin o .*

Proof. Define a distance function on points for each $i \in [1..n]$ as $d_i(q) = \sqrt{|q - p_i|^2 - |o - p_i|^2} = \sqrt{q \cdot q - 2q \cdot p_i}$. This is a form of the power distance defined for circles [3] in which all the circles pass through the origin. The level sets $L_i(\alpha) = \{q \in \mathbb{R}^2 \mid d_i(q) = \alpha\}$ are circles centered at p_i ; for $\alpha > 0$ they are circles that contain the origin o .

Place $c_i(\alpha)$ at the intersection $L_i(\alpha) \cap L_{i+1}(\alpha)$ that lies to the right of $p_i \vec{p}_{i+1}$. We observe that all such $c_i(\alpha)$ lie on the line through the origin perpendicular to $\vec{p}_i \vec{p}_{i+1}$, because this line contains all points q for which $d_i(q)^2 - d_{i+1}(q)^2 = 2q \cdot (p_{i+1} - p_i) = 0$.

If we choose an $\alpha > 0$, therefore, for each $i \in [1..n]$, the trace of $c_i(\alpha)$ goes through the origin, and Lemma 4 says that we can fold to a pyramid with apex projecting to the origin. \square

4 Folding to Bumpy Pyramids with 4 Vertices

We consider the bumpy pyramid, where we allow to fold the base along diagonals that join two vertices $p_i p_j$. To simplify the discussion, we assume that no four points on the base are on the same plane in the resulting polyhedron, so the base is completely triangulated. When $n = 3$, we have the lengths of six edges of a tetrahedron. By Alexandrov's theorem, if the tetrahedron exists, then it is uniquely determined. Its volume can be computed in a constant time using the equation in [10]. Indeed, this could be used instead of Theorem 1 to test foldability when $n = 3$: P folds to a unique tetrahedron iff it has a positive volume computed by the equation. Now we turn to the special case $n = 4$. This case is already not so obvious, and the results in this section are helpful for the general case.

Let $P = (p_1, c_1, p_2, c_2, p_3, c_3, p_4, c_4)$ be a petal polygon. There are two candidates for diagonals to fold,

$p_1 p_3$ or $p_2 p_4$, to make bumpy pyramids from P . For these two candidates, we have the following theorem.

Theorem 6 *For a petal polygon P of 8 points, either (1) no bumpy pyramid can be folded, (2) one convex bumpy pyramid can be folded, or (3) one convex bumpy pyramid and one concave bumpy pyramid can be folded.*

Proof. Suppose that we have a bumpy pyramid by folding along the line $p_2 p_4$; this pyramid consists of two tetrahedra $cp_1 p_2 p_4$ and $cp_3 p_2 p_4$ that share a common triangle $cp_2 p_4$. We can fold it if and only if the three edges, $|p_2 p_4|$, $|cp_2| = \ell_2$, and $|cp_4| = \ell_4$, satisfy the triangle inequality. For each i , let C_i be a circle of radius ℓ_i centered at p_i . The triangle inequality is satisfied if and only if C_2 and C_4 intersect. For the circles C_1 and C_3 , we can use the same argument. Therefore, we have three cases.

Case 1: No pair of (C_1, C_3) and (C_2, C_4) intersect, as in Figure 4(1). In this case, the four triangles are so short that no pyramid can be folded.

Case 2: One of (C_1, C_3) and (C_2, C_4) intersects. Without loss of generality, we assume that C_1 and C_3 intersect as in Figure 4(2). In this case, we cannot make a triangle $cp_2 p_4$ since three edges do not satisfy the triangle inequality. Therefore, only one pyramid can be folded by folding along $p_1 p_3$.

Now we show that the resulting pyramid is convex. To derive a contradiction, we suppose that the pyramid is concave by valley folding along $p_1 p_3$. Consider its intersection with the plane through triangle $p_2 cp_4$: we see that the length of the path from p_2 to p_4 through the apex c is greater than the length of path $p_2 h p_4$ for a point h on $p_1 p_3$. This contradicts that C_2 and C_4 do not intersect. Therefore, to obtain a bumpy pyramid, we must mountain fold along $p_1 p_3$.

Case 3: Both of (C_1, C_3) and (C_2, C_4) intersect, as in Figure 4(3). By Alexandrov's theorem we cannot obtain two convex bumpy pyramids; so we show that one must be convex and the other concave.

We first consider folding along $p_1 p_3$ (Figure 5). The resulting polyhedron can be split into two tetrahedra $cp_1 p_2 p_3$ and $cp_1 p_4 p_3$ by cutting the shared triangle $T = (p_1, p_3, c)$. This T can be depicted twice by joining $p_1 p_3$ and two intersections of C_1 and C_3 (thick lines in Figure 5). Let t_1 and t_2 be the intersection points.

Now we consider two more intersection points; $c_{1,2}$ is of $\overline{c_1(0)c_1(1)}$ and $\overline{c_2(0)c_2(1)}$, and $c_{3,4}$ is of $\overline{c_3(0)c_3(1)}$ and $\overline{c_4(0)c_4(1)}$. Then, by Lemma 4, $c_{1,2}$ and $c_{3,4}$ are both on the line $\overline{t_1 t_2}$. When $c_{1,2}$ is closer to p_2 than $c_{3,4}$ as in the figure, we have to mountain fold along the line $p_1 p_3$ to make the resulting polyhedron. Thus, we assume that $c_{3,4}$ is closer to p_2 . In this case, we have to valley fold along the line $p_1 p_3$ to glue two tetrahedra $cp_1 p_2 p_3$ and $cp_1 p_4 p_3$. However, in this case, when we consider the line $p_2 p_4$, we can conclude that we have to




Figure 4: Three possible cases for bumpy pyramid with $n = 4$: (1) no bumpy pyramid, (2) one convex bumpy pyramid, and (3) two bumpy pyramids are folded, respectively.




Figure 5: (3) Folding along the line p_1p_3 , and its magnification.

mountain fold along the line p_2p_4 to obtain the resulting bumpy pyramid. Therefore, we obtain one convex bumpy pyramid and one concave bumpy pyramid. \square

5 Folding a Convex Bumpy Pyramid

We now turn to the general convex bumpy pyramid problem. A *power diagram* is a partition of the Euclidean plane into polygonal cells defined from a set of circles, where the cell for a given circle C consists of all the points for which the power distance to C is smaller than the power distance to the other circles (see, e.g., [3] for the details). From the cases in the proof of Theorem 6 that created convex bumpy pyramids, we observed that the projections of the c_i as they come together at an apex are sweeping out a power diagram of the vertices of the base B . We first show that this is true in general — that to obtain a convex bumpy pyramid when one exists, we can fold the diagonals of B that are dual to the power diagram edges. Then we show that the

power diagram of the vertices of a convex polygon can be computed in linear time by extending the Voronoi algorithm of Aggarwal *et al.* [1].

Theorem 7 *Let $P = (p_1, c_1, p_2, c_2, \dots, p_n, c_n)$ be a petal polygon with base $B = (p_1, \dots, p_n)$. Let D be the set of diagonals of B that are dual to the power diagram of B with weights $\ell_i = |p_i c_i|$. If P can fold into a convex bumpy pyramid, then it can do so by mountain folding the diagonals of D .*

Proof. The uniqueness of the convex bumpy pyramid follows from Alexandrov’s theorem, so we focus on showing the relation to the power diagram, which is a well-known variation of a Voronoi diagram [3]. For completeness, we include brief sketches of the properties we need; see the literature (e.g., [3]) for more details.

A special form of the power distance was defined for Theorem 5; the general power distance for point p_i with

weight ℓ_i is

$$d_i(q) = \sqrt{|q - p_i|^2 - \ell_i^2}.$$

This is often depicted by drawing the 0-level set, which is a circle C_i centered at p_i of radius ℓ_i .

Use the power distance for the vertices of B to define a Voronoi diagram, which is a decomposition of the plane into maximally connected regions with the same set of closest neighbors. In particular the *edges* $E_{i,j}$ are the non-empty sets defined by pairs of vertices i and j with $i \neq j$ as follows:

$$E_{i,j} = \{q \in \mathbb{R}^2 \mid d_i(q) = d_j(q) \text{ and } \forall k \in [1..n] d_k(q) \leq d_i(q) \longrightarrow k \in \{i, j\}\}.$$

Edges are easily seen to be line segments by expanding $d_i(q) - d_j(q) = 2q(p_i - p_j) + p_i^2 - p_j^2 - (\ell_i^2 - \ell_j^2)$; the line of $E_{i,j}$ contains the intersection points of the circles, $C_i \cap C_j$.

The intersection of halfplanes containing p_i defined by all edges $E_{i,k}$ is the convex cell of p_i , so the edges define a partition of the plane into convex cells, edges, and vertices, which is the *power diagram*. The dual of the power diagram of B tells us which diagonals to use to fold our convex bumpy pyramid: use a diagonal $p_i p_j$ iff $E_{i,j}$ is non-empty.

Because B is convex, we can show that, for all $i \in [1..n]$, power diagram edge $E_{i,i+1}$ goes to infinity. Consider any p_j not an endpoint of the B -edge $\overline{p_i p_{i+1}}$; both vectors $p_i - p_j$ and $p_{i+1} - p_j$ have positive projection on the normal vector $(p_i - p_{i+1})^\perp$, so if we move far enough in that direction, we eventually cross the $E_{i,j}$ and $E_{i+1,j}$ power diagram edges from the p_j side. This completes the characterization of the properties we need from the power diagram.

Now, because the power diagram of vertices of B decomposes the plane into a cell for each p_i , and edges $E_{i,i+1}$ go to infinity, each cell is unbounded and the edges form an acyclic and connected graph. We have a tree, and can consider the infinite edges to be external nodes. A leaf is a tree node that is incident on two external nodes.

This tree gives our folding order. Intuitively, each c_i traces along the edges from corresponding leaf of the power diagram toward the center of the tree. When two leaves meet at their common parent, the corresponding two flaps meet there, and replaced by merged flap. For example, in Figure 6, each c_i starts at a leaf and traces the bold line, which is the power diagram. Then c_4 meets c_3 before c_5 , and c_1 meets c_2 before c_5 . When c_4 meets c_3 , they are glued and two flaps $p_4 p_5 c_4$ and $p_3 p_4 c_3$ are replaced by a new “flap” $p_3 p_5 c'$, where c' is the meeting point of c_3 and c_4 . In other words, we discard the point p_4 from B and proceed recursively with a new P that has a triangle flap $p_3 p_5 c'$ instead of




Figure 6: Power diagram (bold tree) gives the order of gluing the petals.

two flaps $p_4 p_5 c_4$ and $p_3 p_4 c_3$. In the time, all flaps are already partially folded, but no other pair meet with each other.

Then, it is clear that we have to make a mountain fold along $p_3 p_5$ to fold the new flap toward the next meeting point. The recursive folding ends with a base at which all flaps can meet, completing a convex bumpy pyramid.

For each leaf, fold a flap that corresponds to the infinite edge. Suppose that the vertices of P for these flaps are p_{i-1} , c_{i-1} , p_i , c_i , and p_{i+1} . If the traces τ_{i-1} and τ_i are long enough to include the power diagram vertex equidistant to p_{i-1} , p_i , and p_{i+1} , then the flaps for c_{i-1} and c_i will meet at an apex c that projects to the vertex. We then discard the point p_i from B and the leaf from the power diagram, and proceed recursively with a new P that has a triangle flap p_{i-1} , c, p_{i+1} . This flap is already partially folded, but will continue folding forward according to the power diagram edge $E_{i-1, i+1}$, so we may safely fold it back to the plane before continuing.

If each ℓ_i is sufficiently long, it is clear that this procedure certainly fold P into a convex bumpy pyramid. \square

Lemma 8 *The power diagram of B can be computed in linear time.*

Proof. This is a straightforward application of the not-so-straightforward ideas in the Voronoi algorithm of Agarwal *et al.* [1] to the case of the power diagram. Since we have not found it in the literature, we include a detailed sketch.

The basic algorithm is incremental construction, which can be made to run in time proportional to the number of edges of the cell of an inserted vertex in our

case, since we can always know where to begin. First, imagine deleting a site p from a power diagram of B : the neighbors of p 's cell will steal back the area that belonged to the cell. To reinsert p , one can simply walk over the portions of the diagram that will lie in p 's cell, starting from the infinite edge defined by the neighbors of p .

By the way, we could make a randomized incremental construction that runs in two stages: First, randomly delete vertices from B , keeping track of neighbors at the time of deletion. Second, in reverse order, insert vertices incrementally into the power diagram, starting from the infinite edge between the neighbors. This would run in expected linear time, since it is only the point location that necessitates $\Theta(n \log n)$ expected time in randomized incremental construction [5].

For a deterministic algorithm, Aggarwal *et al.* [1] identify a constant fraction of the Voronoi cells that are not adjacent to each other, so that they could be inserted simultaneously in time proportional to the Voronoi size. The same idea works for the power diagram of the vertices of B .

Mark the vertices p_i red or blue to satisfy three rules:

- 1) No two adjacent sites are marked *red*.
- 2) No three adjacent sites are marked *blue*.
- 3) If there is a point in the plane q for which $d_{i-1}(q)$ and $d_{i+1}(q)$ are the two smallest elements of the five element set $\{d_{i-2}(q), d_{i-1}(q), d_i(q), d_{i+1}(q), d_{i+2}(q)\}$ (equivalently, $E_{i-1,i+1}$ is an edge in the power diagram of these five vertices), then p_i is red.

Marking is easy to do in linear time: Initially mark red vertices forced by rule 3, and the rest blue, then while there exists three consecutive blue, change the center to red. The only trouble would be if rule 3 applied to consecutive vertices, but this is not possible due to the structure of power diagrams.

In fact, consecutive red sites cannot have power diagram cells that are adjacent: If cells for p_i and p_j touch and there is a single vertex between them, then that vertex must be red and p_i, p_j are both blue. But there cannot be more than two vertices between consecutive reds, and in such a case the $E_{i,j}$ bisector must run into one of the cells, making the vertex for the other red, so at least one of p_i and p_j are blue.

Aggarwal *et al.* [1] prove a combinatorial lemma that exploits this condition.

Lemma 9 (Aggarwal *et al.* [1]) *Let T be a binary tree embedded in the plane. Each leaf of T has an associated “neighborhood,” which is a connected subtree rooted at that leaf, and leaves adjacent in the topological order around the tree have disjoint neighborhoods. Then there are a fixed fraction of the leaves with disjoint, constant-size neighborhoods, and such leaves can be found in linear time (assuming that neighborhoods can be traced out in breadth-first order).*

Now, we can sketch the divide and conquer of Aggarwal *et al.* [1]: Mark the vertices, and compute the power diagram of the blue by recursion, giving a tree where leaves will be edges that go to infinity between blue vertices that have a red between them. The neighborhood of such a leaf is the region closer to this red vertex than either blue vertex. Lemma 9 says that a constant fraction of the red sites with disjoint, constant-size neighborhoods can be found. These red sites can be merged into the blue power diagram in constant time apiece.

Finally, a constant fraction of the sites remain red; we again compute their power diagram recursively and merge it into the blue power diagram — we can do this in linear total time if we merge connected portions starting and ending with the infinite edges. As Aggarwal *et al.* [1] show, the total time is $O(n)$. \square

With these Lemmas, we prove Theorem 2:

Lemma 10 *If a petal polygon P of $2n$ points folds to a bumpy pyramid, then it folds into a unique convex bumpy pyramid. Moreover, the triangulation of B for folding to the convex bumpy pyramid can be found in linear time using constant precision.*

Proof. By Theorem 7 the power diagram is a tree that gives the folding order, if it is possible. It is easy to follow this tree and perform the folding; since all recursive steps can be phrased in terms of a subpolygon of B with original lengths determining the triangular flaps, the precision does not increase. By Lemma 8, the tree can be computed in linear time and constant precision. \square

6 Bumpy Pyramid of Maximum Volume

We here give two nontrivial examples of petal polygons such that concave ones have larger volumes than convex ones in Figure 7. Using the results by Sabitov [10], we can directly compute their volumes and obtain the claim, but here we give more expressive explanations to reveal the hidden ideas of the polygons. The petal polygon in Figure 7(1) folds to a convex polyhedron of almost zero volume by mountain folding along the line p_1p_3 . On the other hand, it folds to a concave polyhedron of larger volume when we valley fold along the line p_2p_4 (like a fortune cookie). For the other petal polygon in Figure 7(2), when we mountain fold along the line p_2p_4 , we obtain almost a pyramid with a triangular base $p_1p_2p_4$, and the point p_3 is almost on the inner point of the triangle p_2cp_4 . On the other hand, when we valley fold along the line p_1p_3 , although it is skewer and its height is bit reduced, the resulting concave bumpy pyramid has almost a square base, and hence it has a larger volume.




Figure 7: Petal polygons that can fold to form convex or non-convex bumpy pyramids, but the non-convex ones have larger volume than the convex ones.

Therefore, we can conclude that sometimes a non-convex bumpy pyramid has a larger volume than convex one folded from the same petal polygon. This fact motivates us to find a crease pattern of B such that P folds into the bumpy pyramid of maximum volume.

By the following proof, we obtain Theorem 3, as stated in Introduction.

Proof. (of Theorem 3) We use dynamic programming in which each subproblem $S(i, k)$ is a sequence of $1 \leq k < n$ consecutive edges of B , using $k + 1$ vertices p_i to p_{i+k} . The weight of a subproblem, $w(i, k)$, is the maximum volume of the bumpy pyramid folded from the petal polygon with vertices $p_i, c_i, \dots, c_{i+k-1}, p_{i+k}, c'$, where c' is the point with lengths $|c'p_i| = \ell_i$ and $|c'p_{i+1}| = \ell_{i+1}$. Thus, when $k = 1$ we have two flaps that fold together, $w(i, 1) = 0$, and when $k = n - 1$, we have P back again and $w(1, n - 1)$ is the maximum volume that we seek.

We say that (i, j, k) , with $1 \leq i \leq n, 1 \leq j < k < n$ is *valid* if a six-vertex petal polygon with base $\Delta p_i p_{i+j} p_{i+k}$ and flaps with corresponding lengths ℓ_i, ℓ_{i+j} and ℓ_{i+k} can be folded to a tetrahedron; let $V(i, j, k)$ denote the volume of this tetrahedron. For $1 < k < n$,

$$w(i, k) = \begin{cases} \max_{1 \leq j < k} V(i, j, k) + w(i, j) + w(i + j, k) & \text{if } (i, j, k) \text{ is valid,} \\ -\infty & \text{otherwise.} \end{cases}$$

Thus, we have n^2 subproblems, and the computation for each can be performed in $O(n)$ time if we assume

a Real RAM model. By storing with each subproblem the index j at which the maximum occurred, we can recover the crease pattern that gives the maximum. \square

We note that the dynamic programming in the proof of Theorem 3 can be modified to count the number of valid bumpy pyramids. Therefore we have the following corollary:

Corollary 11 *Let P be a petal polygon whose base B has n vertices. The number of the valid bumpy pyramids folded from P (by triangulations of B) can be computed in $O(n^3)$ time.*

7 Concluding Remarks

In this paper, we consider polyhedra folded from a petal polygon P . Although the complete characterization of the property is still open, we give the first nontrivial steps to the new problem. Especially, we give the first nontrivial and efficient algorithms for Alexandrov's theorem of finding a unique crease pattern if it exists, by restricting our attention to the special case of petal polygons.

We close by mentioning the potential relationship between this problem and the big open problem that asks whether every convex polyhedron has a non-overlapping net. When we attempt to unfold a general triangulated convex polyhedron onto the plane, one natural way is to first cut and open at a vertex of the convex polyhedron. In the case, we cut all edges incident to the vertex, and open the triangle flaps (as long as faces are triangulated). After that, the natural reduction process would remove the triangle flaps and attach a bumpy (triangulated) face along the open area. This process can be regarded as removing a bumpy pyramid from the polyhedron to reduce the number of vertices. If we have a complete characterization of bumpy pyramids, it may give us the first step of the induction.

The last author thanks Teruo Nishiyama, who introduced the petal pyramid folding problem for $n = 4$. His question led us to this rich vein of problems about bumpy pyramids.

References

- [1] A. Aggarwal, L. J. Guibas, J. Saxe, and P. W. Shor. A Linear-Time Algorithm for Computing the Voronoi Diagram of a Convex Polygon. *Discrete Computational Geometry*, 4(1):591–604, 1989.
- [2] B. Aronov and J. O'Rourke. Nonoverlap of the Star Unfolding. *Discrete Computational Geometry*, 8:219–250, 1992.
- [3] F. Aurenhammer. Power Diagrams: Properties, Algorithms and Applications. *SIAM J. Comput.*, 16:78–96, 1987.

-
- [4] A. I. Bobenko and I. Izestiev. Alexandrov’s theorem, weighted Delaunay triangulations, and mixed volumes. arXiv:math.DG/0609447, February 2008.
- [5] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 3 edition, 2008.
- [6] E. D. Demaine and J. O’Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, 2007.
- [7] A. Dürer. *Underweysung der messung, mit den zirckel un richtscheyt, in Linien ebenen unnd gantzen corporen*. 1525.
- [8] D. Kane, G. N. Price, and E. D. Demaine. A pseudopolynomial algorithm for Alexandrov’s Theorem. In *11th Algorithms and Data Structures Symposium (WADS 2009)*, pages 435–446. Lecture Notes in Computer Science Vol. 5664, Springer-Verlag, 2009.
- [9] J. O’Rourke. *How to Fold It: The Mathematics of Linkage, Origami and Polyhedra*. Cambridge University Press, 2011.
- [10] I. K. Sabitov. The Volume as a Metric Invariant of Polyhedra. *Discrete and Computational Geometry*, 20:405–425, 1998.