Voronoi Cells

Bernd Sturmfels

February 7, 2023

Every real algebraic variety determines a Voronoi decomposition of its ambient Euclidean space. Each Voronoi cell is a convex semialgebraic set in the normal space of the variety at a point. We here study these Voronoi cells, with primary focus on their algebraic boundaries.

We begin with the familiar case when X is a finite subset of the Euclidean space $\mathbb{R}^n$. The *Voronoi cell* of a point $y \in X$ consists of all points whose closest point in X is y, i.e.

$$\mathrm{Vor}_X(y) \;\; := \;\; \big\{\, u \in \mathbb{R}^n : y \in \operatorname*{arg\,min}_{x \in X} \|x-u\|^2 \,\big\}. \tag{1}$$

This is a convex polyhedron with at most $|X|-1$ facets. The study of these cells, and how they depend on X, is ubiquitous in computational geometry and its numerous applications.

Proposition 1. *The Voronoi cell of a point y in the finite set $X \subset \mathbb{R}^n$ is the polyhedron*

$$\mathrm{Vor}_X(y) \;=\; \big\{\, u \in \mathbb{R}^n \,:\, u \cdot (x-y) \,\leq\, \frac{1}{2}\big(||x||^2 - ||y||^2\big) \;\textit{for all}\; x \in X\backslash\{y\}\big\}. \tag{2}$$

Proof. By definition, $\mathrm{Vor}_X(y)$ consists of all points u such that $||x-u||^2 - ||y-u||^2$ is nonnegative for all $x \in X\backslash\{y\}$. But, this expression is equal to $||x||^2 - ||y^2|| - 2\,u\cdot(x-y)$. The main point is that the quadratic term drops out, so the expression is linear in u. □

The collection of all Voronoi cells, as y ranges over X, is known as the *Voronoi diagram* of X. The Voronoi diagram is a polyhedral subdivision of $\mathbb{R}^n$ into finitely many convex cells.

We now shift gears, and we replace the finite set X by a real algebraic variety of positive dimension. As before, the ambient space is $\mathbb{R}^n$ with its standard Euclidean metric. We seek the Voronoi diagram $\{\mathrm{Vor}_X(y)\}_{y\in X}$ in $\mathbb{R}^n$ where y runs over all (infinitely many) points in X.

One approach is to take a large but finite sample from X and to consider the Voronoi diagram of that sample. This is a finite approximation to the desired limit object. By taking finer and finer samples, the Voronoi diagram should converge nicely to a subdivision with infinitely many regions. The Voronoi cells in the limit are convex sets. However, for $n \geq 3$, they are generally not polyhedra. This process was studied by Brandt and Weinstein in [3] for the case when $n = 2$ and X is a curve. In [3, Figure 1] we see this for a quartic curve. The authors posted a delightful YouTube video, called *Mathemaddies' Ice Cream Map*. Please do watch that movie! Their curve X is the shoreline that separates the city of Berkeley from the San Francisco Bay. One hopes to find many ice cream shops at the shore.

Let X be a real algebraic variety of codimension c in $\mathbb{R}^n$, and $y \in X$. The Voronoi cell $\mathrm{Vor}_X(y)$ is defined as before. It consists of all points u in $\mathbb{R}^n$ such that y is closer or equal to u than any other point $x \in X$. The equation (2) still holds, so $\mathrm{Vor}_X(y)$ is a convex set.

Proposition 2. *Suppose that y is a smooth point in X. The Voronoi cell* $\mathrm{Vor}_X(y)$ *is a convex semialgebraic set of dimension c. It is contained in the* normal space

$$N_X(y) \;=\; \bigl\{u \in \mathbb{R}^n \,:\, u-y \;\textit{is perpendicular to the tangent space of}\; X \;\textit{at}\; y\bigr\} \;\simeq\; \mathbb{R}^c.$$

Proof. Fix $u \in \mathrm{Vor}_X(y)$. Consider any point x in X that is close to y, and set $v = x-y$. The inequality in (2) implies $u\cdot v \,\leq\, \frac{1}{2}(||y+v||^2 - ||y||^2) \,=\, y\cdot v + \frac{1}{2}||v||^2$. For any w in the tangent space of X at y, there exists $v = \epsilon w + O(\epsilon^2)$ such that $x = y+v$ is in X. The inequality above yields $u\cdot w \leq y\cdot w$, and the same with $-w$ instead of w. Then $(u-y)\cdot w = 0$, and hence u is in the normal space $N_X(y)$. We already argued that $\mathrm{Vor}_X(y)$ is convex. It is semialgebraic, by Tarski's Theorem on Quantifier Elimination, which allows us to eliminate x from the formula (2). Finally, the Voronoi cell $\mathrm{Vor}_X(y)$ is full-dimensional in the c-dimensional space $N_X(y)$ because every point u in an ϵ-neighborhood of y has a unique closest point in X. If $u \in N_X(y)$ then that closest point must be y, by the same inequality as above. □

The topological boundary of $\mathrm{Vor}_X(y)$ in $N_X(y)$ is denoted by $\partial\,\mathrm{Vor}_X(y)$. It consists of points in $N_X(y)$ that have at least two closest points in X, including y. We are interested in the *algebraic boundary* $\partial_{\mathrm{alg}}\mathrm{Vor}_X(y)$. This is the hypersurface in the complex affine space $N_X(y)_{\mathbb{C}} \simeq \mathbb{C}^c$ obtained as the Zariski closure of $\partial\,\mathrm{Vor}_X(y)$ over the field of definition of X. The degree of this hypersurface is denoted $\delta_X(y)$ and called the *Voronoi degree* of X at y. If X is irreducible and y is a general point on X, then this degree does not depend on y.

Example 3 (Surfaces in 3-space)**.** Fix a general inhomogeneous polynomial $f \in \mathbb{Q}[x_1, x_2, x_3]$ of degree $d \geq 2$ and let $X = V(f)$ be its surface in $\mathbb{R}^3$. The normal space at a general point $y \in X$ is the line $N_X(y) = \{y + \lambda(\nabla f)(y) \,:\, \lambda \in \mathbb{R}\}$. The Voronoi cell $\mathrm{Vor}_X(y)$ is a (possibly unbounded) line segment in $N_X(y)$ that contains the point y. The boundary $\partial\,\mathrm{Vor}_X(y)$ consists of at most two points from among the zeros of an irreducible polynomial in $\mathbb{Q}[\lambda]$. We shall see that this univariate polynomial has degree $d^3 + d - 7$. Its complex zeros form the algebraic boundary $\partial_{\mathrm{alg}}\mathrm{Vor}_X(y)$. Thus, the Voronoi degree of the surface X is d^3+d-7.

Note that, in this example, our hypothesis "over the field of definition" becomes important. The $\mathbb{Q}$-Zariski closure of one boundary point is the collection of all $d^3 + d - 7$ points.

For a numerical example, let $d = 2$ and fix $y = (0, 0, 0)$ and $f = x_1^2 + x_2^2 + x_3^2 - 3x_1x_2 - 5x_1x_3 - 7x_2x_3 + x_1 + x_2 + x_3$. Let $r_0 \approx -0.209$, $r_1 \approx -0.107$, $r_2 \approx 0.122$ be the roots of the cubic polynomial $368\lambda^3 + 71\lambda^2 - 6\lambda - 1$. The Voronoi cell $\mathrm{Vor}_X(y)$ is the line segment connecting the points (r_1, r_1, r_1) and (r_2, r_2, r_2). The topological boundary $\partial\,\mathrm{Vor}_X(y)$ consists of $(r_1, r_1, r_1), (r_2, r_2, r_2)$, whereas the algebraic boundary $\partial_{\mathrm{alg}}\,\mathrm{Vor}_X(y)$ also contains (r_0, r_0, r_0).

The cubic polynomial in λ was found with the algebraic method that is described below. Namely, the Voronoi ideal in (3) equals $\mathrm{Vor}_I(0) = \bigl\langle u_1 - u_3,\, u_2 - u_3,\, 368u_3^3 + 71u_3^2 - 6u_3 - 1\bigr\rangle$. This is a maximal ideal in $\mathbb{Q}[u_1, u_2, u_3]$, which defines a field extension of degree 3 over $\mathbb{Q}$.

Example 4 (Curves in 3-space)**.** Let X be a general algebraic curve in $\mathbb{R}^3$. For $y \in X$, the Voronoi cell $\mathrm{Vor}_X(y)$ is a convex set in the normal plane $N_X(y) \simeq \mathbb{R}^2$. Its algebraic boundary $\partial_{\mathrm{alg}}\mathrm{Vor}_X(y)$ is a plane curve of degree $\delta_X(y)$. This Voronoi degree can be expressed in terms of the degree and genus of X. Specifically, if X is the intersection of two general quadrics in $\mathbb{R}^3$, then the Voronoi degree is 12. Figure 1 shows one such quartic space curve X together

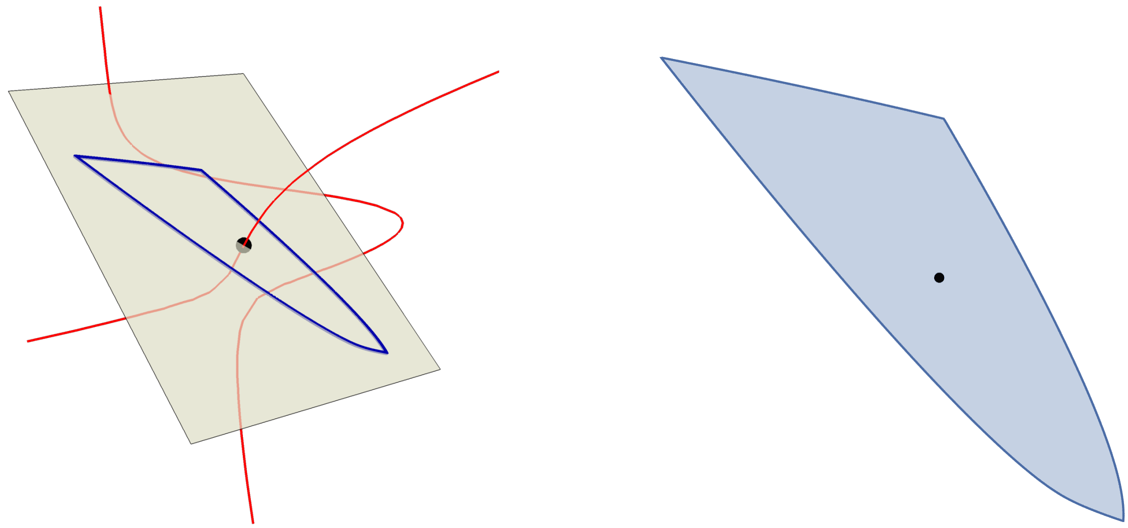

Figure 1: A quartic space curve, shown with the Voronoi cell in one of its normal planes.

with the normal plane at a point $y \in X$. The Voronoi cell $\mathrm{Vor}_X(y)$ is the planar convex region highlighted on the right. Its algebraic boundary $\partial_{\mathrm{alg}}\mathrm{Vor}_X(y)$ is a curve of degree $\delta_X(y) = 12$. The topological boundary $\partial\,\mathrm{Vor}_X(y)$ is only a very small subset of that algebraic boundary.

Metric algebraic geometry is concerned with properties of real algebraic varieties that depend on a distance metric. Key concepts include the Euclidean distance degree [8], distance function [14], bottlenecks [7, 10], reach, offset hypersurfaces, medial axis [13], and cut locus [5]. Voronoi cells are also an important topic in metric algebraic geometry. We here consider them only for the Euclidean metric, but it also makes much sense to study Voronoi cells with respect to Kullback-Leibler divergence [1] or Wasserstein distance [2].

We study the Voronoi decomposition to answer the question for any point in ambient space, “What point on the variety X am I closest to?” Another question one might ask is, “How far do we have to get away from X before there is more than one answer to the closest point question?” The union of the boundaries of the Voronoi cells is the locus of points in $\mathbb{R}^n$ that have more than one closest point on X. This set is called the *medial axis* (or *cut locus*) of the variety. The distance from the variety to its medial axis, which is the answer to the “how far” question, is called the *reach* of X. This quantity is of interest, for example, in topological data analysis, as it is the main quantity determining the density of sample points needed to compute the persistent homology of X. We refer to [4, 9] for studies on sampling at the interface of topological data analysis with metric algebraic geometry. The distance from a point y on X to the variety's medial axis could be considered the *local reach* of X. Equivalently, this is the distance from y to the boundary of its Voronoi cell $\mathrm{Vor}_X(y)$.

The material that follows is based on the article [6]. We begin with the exact symbolic computation of the Voronoi boundary at y from the equations that define X. This uses a Gröbner-based algorithm whose input is y and the ideal of X and whose output is the ideal defining $\partial_{\mathrm{alg}}\mathrm{Vor}_X(y)$. This is followed by formulas for the Voronoi degree $\delta_X(y)$ when X, y are sufficiently general and $\dim(X) \leq 2$. Thereafter we study the case when y is a low rank matrix and X is the variety of these matrices. This relies on the *Eckart-Young Theorem*.

We now describe Gröbner basis methods for finding the Voronoi boundaries of a given variety. We start with an ideal $I = \langle f_1, f_2, \ldots, f_m \rangle$ in $\mathbb{Q}[x_1, \ldots, x_n]$ whose real variety $X = V(I) \subset \mathbb{R}^n$ is assumed to be nonempty. We assume that I is real radical and prime, so that $X_{\mathbb{C}}$ is an irreducible variety in $\mathbb{C}^n$ whose real points are Zariski dense. Our aim is to compute the Voronoi boundary of a given point $y \in X$. In our examples, the coordinates of the point y and the coefficients of the polynomials f_i are rational numbers. Under these assumptions, the following computations can be done in polynomial rings over $\mathbb{Q}$.

Fix the polynomial ring $R = \mathbb{Q}[x_1, \ldots, x_n, u_1, \ldots, u_n]$ where $u = (u_1, \ldots, u_n)$ is an additional unknown point. The *augmented Jacobian* of X at x is the following matrix of size $(m+1) \times n$ with entries in R. It contains the n partial derivatives of the m generators of I:

$$J_I(x, u) \quad := \quad \begin{bmatrix} u - x \\ (\nabla f_1)(x) \\ \vdots \\ (\nabla f_m)(x) \end{bmatrix}$$

Let N_I denote the ideal in R generated by I and the $(c+1) \times (c+1)$ minors of the augmented Jacobian $J_I(x, u)$, where c is the codimension of the given variety $X \subset \mathbb{R}^n$. The ideal N_I in R defines a subvariety of dimension n in $\mathbb{R}^{2n}$, namely the *Euclidean normal bundle* of X. Its points are pairs (x, u) where x is a point in X and u lies in the normal space of X at x.

Example 5 (Cuspidal cubic)**.** Let $n = 2$ and $I = \langle\, x_1^3 - x_2^2 \,\rangle$, so $X = V(I) \subset \mathbb{R}^2$ is a cubic curve with a cusp at the origin. The ideal of the Euclidean normal bundle of X is

$$N_I \;=\; \Big\langle\, x_1^3 - x_2^2 \,,\, \det\begin{pmatrix} u_1 - x_1 & u_2 - x_2 \\ 3x_1^2 & -2x_2 \end{pmatrix} \Big\rangle.$$

For $y \in \mathbb{R}^2$, let $N_I(y)$ denote the linear ideal that is obtained from N_I by replacing the unknown point x by the given point y. For instance, for $y = (4, 8)$ we obtain $N_I(y) = \langle u_1 + 3u_2 - 28 \rangle$. We now define the *critical ideal* of the variety X at the point y as

$$C_I(y) \quad = \quad N_I \,+\, N_I(y) \,+\, \langle\, \|x - u\|^2 - \|y - u\|^2 \rangle \quad \subset \quad R.$$

The variety of $C_I(y)$ consists of pairs (u, x) such that x and y are equidistant from u and both are critical points of the distance function from u to X. The *Voronoi ideal* is the following ideal in $\mathbb{Q}[u_1, \ldots, u_n]$. It is obtained from the critical ideal by saturation and elimination:

$$\mathrm{Vor}_I(y) \quad = \quad \big(\, C_I(y) \,:\, \langle x - y \rangle^\infty \big) \,\cap\, \mathbb{Q}[u_1, \ldots, u_n]. \tag{3}$$

The geometric interpretation of each step in our construction implies the following result:

Proposition 6. *The affine variety in $\mathbb{C}^n$ defined by the Voronoi ideal $\mathrm{Vor}_I(y)$ contains the algebraic Voronoi boundary $\partial_{\mathrm{alg}} \mathrm{Vor}_X(y)$ of the given real variety X at its point y.*

Remark 7. The verb "contains" sounds a bit weak, but it is stronger than it may seem. In generic situations, the ideal $\mathrm{Vor}_I(y)$ will be prime, and it defines an irreducible hypersurface in the normal space $N_I(y)$. This hypersurface equals the algebraic Voronoi boundary, so containment is an equality. We saw this in Example 3. For special data, $\mathrm{Vor}_I(y)$ usually defines a hypersurface in $N_I(y)$, but it can have extraneous components, which we remove.

Example 8. For the point $y = (4,8)$ on the cuspidal cubic X in Example 5, we have $N_I(y) = \langle u_1 + 3u_2 - 28\rangle$. Going through the steps above, we find that the Voronoi ideal is

$$\mathrm{Vor}_I(y) \;=\; \langle u_1 - 28, u_2\rangle \;\cap\; \langle u_1 + 26, u_2 - 18\rangle \;\cap\; \langle u_1 + 3u_2 - 28,\, 27u_2^2 - 486u_2 + 2197\rangle.$$

The third component has no real roots and is hence extraneous. The Voronoi boundary consists of two points: $\partial\,\mathrm{Vor}_X(y) = \{(28,0),(-26,18)\}$. The Voronoi cell $\mathrm{Vor}_X(y)$ is the line segment connecting these points. This segment is shown in green in Figure 2. Its right endpoint $(28,0)$ is equidistant from y and the point $(4,-8)$. Its left endpoint $(-26,18)$ is equidistant from y and the point $(0,0)$, whose Voronoi cell is discussed in Remark 9.

The cuspidal cubic X is very special. If we replace X by a general cubic (defined over $\mathbb{Q}$) in the affine plane, then $\mathrm{Vor}_I(y)$ is generated modulo $N_I(y)$ by an irreducible polynomial of degree eight in $\mathbb{Q}[u_2]$. Thus, the expected Voronoi degree of (affine) plane cubics is $\delta_X(y) = 8$.

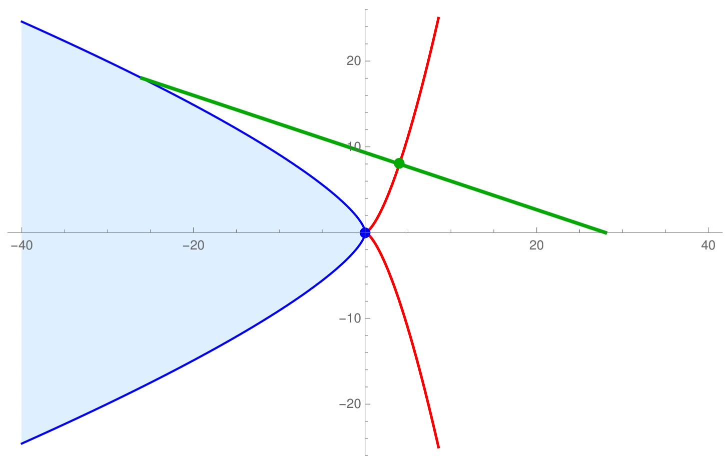

Figure 2: The cuspidal cubic is shown in red. The Voronoi cell of a smooth point is a green line segment. The Voronoi cell of the cusp is the convex region bounded by the blue curve.

Remark 9 (Singularities)**.** Voronoi cells at singular points can be computed with the same procedure as above. However, these Voronoi cells generally have higher dimensions. For an illustration, consider the cuspidal cubic, and let $y = (0,0)$ be the cusp. A Gröbner basis computation yields the Voronoi boundary $27u_2^4 + 128u_1^3 + 72u_1u_2^2 + 32u_1^2 + u_2^2 + 2u_1$. The Voronoi cell is the two-dimensional convex region bounded by this quartic, shown in blue in Figure 2. The Voronoi cell might also be empty at a singularity. This happens for instance for $V(x_1^3 + x_1^2 - x_2^2)$, which has an ordinary double point at $y=(0,0)$. In general, the cell dimension depends on both the embedding dimension and the branches of the singularity.

Proposition 6 gives an algorithm for computing the Voronoi ideal $\mathrm{Vor}_I(y)$ when y is a smooth point in $X = V(I)$. Experiments with `Macaulay2` [12] are reported in [6]. For small enough instances, the computation terminates and we obtain the defining polynomial of the Voronoi boundary $\partial_{\mathrm{alg}}\mathrm{Vor}_X(y)$. This polynomial is unique modulo the linear ideal of the normal space $N_I(y)$. For larger instances, we can only compute the degree of $\partial_{\mathrm{alg}}\mathrm{Vor}_X(y)$

but not its equation. This is done by working over a finite field and adding $c-1$ random linear equations in $u_1, \ldots, u_n$ in order to get a zero-dimensional polynomial system.

Computations are easiest to set up for the case of hypersurfaces ($c = 1$). We explore random polynomials f of degree d in $\mathbb{Q}[x_1, \ldots, x_n]$, both inhomogeneous and homogeneous. These are chosen among those that vanish at a preselected point y in $\mathbb{Q}^n$. In each iteration, the Voronoi ideal $\mathrm{Vor}_I(y)$ from (3) was found to be zero-dimensional. In fact, $\mathrm{Vor}_I(y)$ is a maximal ideal in $\mathbb{Q}[u_1, \ldots, u_n]$, and $\delta_X(y)$ is the degree of the associated field extension.

We summarize our results in Tables 1 and 2, and we extract conjectural formulas.

$n\backslash d$	2	3	4	5	6	7	8	$\delta_X(y) = \mathrm{degree}(\mathrm{Vor}_{\langle f\rangle}(y))$
1	1	2	3	4	5	6	7	$d-1$
2	2	8	16	26	38	52	68	d^2+d-4
3	3	23	61	123	215	343		d^3+d-7
4	4	56	202	520	1112			$d^4-d^3+d^2+d-10$
5	5	125	631					$d^5-2d^4+2d^3+d-13$
6	6	266	1924					$d^6-3d^5+4d^4-2d^3+d^2+d-16$
7	7	551						$d^7-4d^6+7d^5-6d^4+3d^3+d-19$

Table 1: The Voronoi degree of an inhomogeneous polynomial f of degree d in $\mathbb{R}^n$.

$n\backslash d$	2	3	4	5	6	7	8	$\delta_X(y) = \mathrm{degree}(\mathrm{Vor}_{\langle f\rangle}(y))$
2	2	4	6	8	10	12	14	$2d-2$
3	3	13	27	45	67	93	123	$2d^2-5$
4	4	34	96	202				$2d^3-2d^2+2d-8$
5	5	79	309					$2d^4-4d^3+4d^2-11$
6	6	172						$2d^5-6d^4+8d^3-4d^2+2d-14$
7	7	361						$2d^6-8d^5+14d^4-12d^3+6d^2-17$

Table 2: The Voronoi degree of a homogeneous polynomial f of degree d in $\mathbb{R}^n$.

Conjecture 10. The Voronoi degree of a generic hypersurface of degree d in $\mathbb{R}^n$ equals

$$(d-1)^n + 3(d-1)^{n-1} + \tfrac{4}{d-2}((d-1)^{n-1} - 1) - 3n.$$

The Voronoi degree of the cone of a generic homogeneous polynomial of degree d in $\mathbb{R}^n$ is

$$2(d-1)^{n-1} + \tfrac{4}{d-2}((d-1)^{n-1} - 1) - 3n + 2.$$

Both parts of this conjecture are proved for $n \leq 3$ in [6, Section 4], where the geometric theory of Voronoi degrees of low-dimensional varieties is developed. The case $d = 2$ was analyzed in [5, Proposition 5.8]. In general, for $n \geq 4$ and $d \geq 3$, the problem is open.

To recap, the algebraic boundary of the Voronoi cell $\mathrm{Vor}_X(y)$ is a hypersurface in the normal space to a variety $X \subset \mathbb{R}^n$ at a point $y \in X$. We shall present formulas for the degree $\delta_X(y)$ of that hypersurface when X is a curve or a surface. All proofs appear in [6, Section 6].

We identify X and $\partial_{\mathrm{alg}}\mathrm{Vor}_X(y)$ with their Zariski closures in complex projective space $\mathbb{P}^n$, so there is a natural assigned hyperplane at infinity. We say that X is in *general position* in $\mathbb{P}^n$ if the hyperplane at infinity intersects X transversally, i.e. that the intersection is smooth.

Theorem 11. *Let $X \subset \mathbb{P}^n$ be a curve of degree d and geometric genus g with at most ordinary multiple points as singularities. The Voronoi degree at a general point $y \in X$ equals*

$$\delta_X(y) \;=\; 4d + 2g - 6,$$

provided X is in general position in $\mathbb{P}^n$.

Example 12. If X is a smooth curve of degree d in the plane, then $2g - 2 = d(d-3)$, so

$$\delta_X(y) \;=\; d^2 + d - 4.$$

This confirms our experimental results in the row $n = 2$ of Table 1.

Example 13. If X is a rational curve of degree d, then $g = 0$ and hence $\delta_X(y) = 4d - 6$. If X is an elliptic curve, so the genus is $g = 1$, then we have $\delta_X(y) = 4d - 4$. A space curve with $d = 4$ and $g = 1$ was studied in Example 4. Its Voronoi degree equals $\delta_X(y) = 12$.

Theorem 11 is [6, Theorem 5.1]. The general position assumption is essential. For an example, let X be the twisted cubic curve in $\mathbb{P}^3$, with affine parameterization $t \mapsto (t, t^2, t^3)$. Here $g = 0$ and $d = 3$, so the expected Voronoi degree is 6. However, a computation shows that $\delta_X(y) = 4$. This drop is explained by the fact that the plane at infinity in $\mathbb{P}^3$ intersects the curve X in a triple point. After a general linear change of coordinates in $\mathbb{P}^3$, which amounts to a linear fractional transformation in $\mathbb{R}^3$, we correctly find $\delta_X(y) = 6$.

We next present a formula for the Voronoi degree of a surface X which is smooth and irreducible in $\mathbb{P}^n$. Our formula is in terms of its degree d and two further invariants. The first, denoted $\chi(X) := c_2(X)$, is the topological Euler characteristic. This is equal to the degree of the second Chern class of the tangent bundle. The second invariant, denoted $g(X)$, is the genus of the curve obtained by intersecting X with a general smooth quadratic hypersurface in $\mathbb{P}^n$. Thus, $g(X)$ is the quadratic analogue to the usual sectional genus of the surface X.

Theorem 14 (Theorem 5.4 in [6])**.** *Let $X \subset \mathbb{P}^n$ be a smooth surface of degree d. Then*

$$\delta_X(y) \;=\; 3d + \chi(X) + 4g(X) - 11,$$

provided the surface X is in general position in $\mathbb{P}^n$ and y is a general point on X.

Example 15. If X is a smooth surface in $\mathbb{P}^3$ of degree d, then $\chi(X) = d(d^2 - 4d + 6)$, by [11, Ex 3.2.12]. A smooth quadratic hypersurface section of X is an irreducible curve of degree (d, d) in $\mathbb{P}^1 \times \mathbb{P}^1$. The genus of such a curve is $g(X) = (d-1)^2$. We conclude that

$$\delta_X(y) \;=\; 3d + d(d^2 - 4d + 6) \,+\, 4(d-1)^2 \,-\, 11 \;=\; d^3 + d - 7.$$

This confirms our experimental results in the row $n = 3$ of Table 1.

Example 16. Let X be the Veronese surface of order e in $\mathbb{P}^{\binom{e+1}{2}-1}$, taken after a general linear change of coordinates in that ambient space. The degree of X equals $d = e^2$. We have $\chi(X) = \chi(\mathbb{P}^2) = 3$, and the general quadratic hypersurface section of X is a curve of genus $g(X) = \binom{2e-1}{2}$. We conclude that the Voronoi degree of X at a general point y equals

$$\delta_X(y) \;=\; 3e^2 \,+\, 3 \,+\, 2(2e{-}1)(2e{-}2) \,-\, 11 \;=\; 11e^2 - 12e - 4.$$

For instance, for the quadratic Veronese surface in $\mathbb{P}^5$ we have $e = 2$ and hence $\delta_X(y) = 16$. This is smaller than the number 18 found in Example 22, since back then we were dealing with the cone over the Veronese surface in $\mathbb{R}^6$, and not with the Veronese surface in $\mathbb{R}^5 \subset \mathbb{P}^5$.

We finally consider affine surfaces defined by homogeneous polynomials. Namely, let $X \subset \mathbb{R}^n$ be the affine cone over a general smooth curve of degree d and genus g in $\mathbb{P}^{n-1}$.

Theorem 17 (Theorem 5.7 in [6])**.** *If $X \subset \mathbb{R}^n$ is the cone over a smooth curve in $\mathbb{P}^{n-1}$ then*

$$\delta_X(y) \;=\; 6d + 4g - 9,$$

provided that the curve is in general position and y is a general point.

Example 18. If $X \subset \mathbb{R}^3$ is the cone over a smooth curve of degree d in $\mathbb{P}^2$, then $2g - 2 = d(d-3)$. Hence the Voronoi degree of X is

$$\delta_X(y) \;=\; 2d^2 - 5.$$

This confirms our experimental results in the row $n = 3$ of Table 2.

To conclude, we comment on the assumptions made in our theorems. We assumed that the variety X is in general position in $\mathbb{P}^n$. If this is not satisfied, then the Voronoi degree may drop. The point here is that the Voronoi ideal $\mathrm{Vor}_I(y)$ depends polynomially on the description of X, and the degree of this zero-dimensional ideal can only go down – and not up – when that description specializes. Making this statement precise would require a technical discussion of families in algebraic geometry, a topic best left to the experts on foundations. Nonetheless, the technique introduced in the next section can be adapted to determine the correct value. As an illustration, we consider the affine Veronese surface (Example 16).

Example 19. Let $X \subset \mathbb{P}^5$ be the Veronese surface with affine parametrization $(s,t) \mapsto (s,t,s^2,st,t^2)$. The hyperplane at infinity intersects X in a double conic, so X is not in general position. In the next section, we will show that the true Voronoi degree is $\delta_X(y) = 10$. For the Frobenius norm, the Voronoi degree drops further. For this, we shall derive $\delta_X(y) = 4$.

We now turn to the case of great interest in applications. Let X be the variety of real $m \times n$ matrices of rank $\leq r$. We consider two natural norms on the space $\mathbb{R}^{m\times n}$ of real $m \times n$ matrices. First, we have the *Frobenius norm* $\|U\|_F := \sqrt{\sum_{ij} U_{ij}^2}$. And second, we have the *spectral norm* $\|U\|_2 := \max_i \sigma_i(U)$ which extracts the largest singular value of the matrix U.

Fix a rank r matrix V in X. This is a nonsingular point in X. We consider the Voronoi cell $\mathrm{Vor}_X(V)$ with respect to the Frobenius norm. This is consistent with our setting because the Frobenius norm agrees with the Euclidean norm on $\mathbb{R}^{m\times n}$. This identification will no longer be valid when we restrict to the subspace of symmetric matrices.

Let $U \in \mathrm{Vor}_X(V)$, i.e. the closest point to U in the rank r variety X is the matrix V. By the Eckart-Young Theorem, the matrix V is derived from U by computing the singular value decomposition $U = \Sigma_1 D \Sigma_2$. Here Σ_1 and Σ_2 are orthogonal matrices of size $m \times m$ and $n \times n$ respectively, and D is a nonnegative diagonal matrix whose entries are the singular values. Let $D^{[r]}$ be the matrix that is obtained from D by replacing all singular values except for the r largest ones by zero. Then, according to Eckart-Young, we have $V = \Sigma_1 \cdot D^{[r]} \cdot \Sigma_2$.

Remark 20. The Eckart-Young Theorem works for both the Frobenius norm and the spectral norm. This means that $\mathrm{Vor}_X(V)$ is also the Voronoi cell for the spectral norm.

The following theorem describes the Voronoi cells for low-rank matrix approximation.

Theorem 21. *Let V be an $m \times n$-matrix of rank r. The Voronoi cell $\mathrm{Vor}_X(V)$ is congruent up to scaling to the unit ball in the spectral norm on the space of $(m-r)\times(n-r)$-matrices.*

Before we present the proof, let us first see why the statement makes sense. The determinantal variety X has dimension $rm + rn - r^2$ in an ambient space of dimension mn. The dimension of the normal space at a point is the difference of these two numbers, so it equals $(m-r)(n-r)$. Every Voronoi cell is a full-dimensional convex body in the normal space. Next consider the case $m = n$ and restrict to the space of diagonal matrices. Now X is the set of vectors in $\mathbb{R}^n$ having at most r nonzero coordinates. This is a reducible variety with $\binom{n}{r}$ components, each a coordinate subspace. For a general point y in such a subspace, the Voronoi cell $\mathrm{Vor}_X(y)$ is a convex polytope. It is congruent to a regular cube of dimension $n-r$, which is the unit ball in the L^∞-norm on $\mathbb{R}^{n-r}$. Theorem 21 describes the orbit of this picture under the action of the two orthogonal groups on $\mathbb{R}^{m\times n}$. For example, consider the special case $n = 3, r = 1$. Here, X consists of the three coordinate axes in $\mathbb{R}^3$. The Voronoi decomposition of this curve decomposes $\mathbb{R}^3$ into squares, each normal to a different point on the three lines. The image of this picture under orthogonal transformations is the Voronoi decomposition of $\mathbb{R}^{3\times 3}$ associated with the affine variety of rank 1 matrices. That variety has dimension 5, and each Voronoi cell is a 4-dimensional convex body in the normal space.

Proof of Theorem 21. The Voronoi cell is invariant under orthogonal transformations. We may therefore assume that the matrix $V = (v_{ij})$ satisfies $v_{11} \geq v_{22} \geq \cdots \geq v_{rr} = u > 0$ and $v_{ij} = 0$ for all other entries. The Voronoi cell of the diagonal matrix V consists of matrices U whose block-decomposition into $r + (m-r)$ rows and $r + (n-r)$ columns satisfies

$$\begin{pmatrix} I & 0 \\ 0 & T_1 \end{pmatrix} \cdot \begin{pmatrix} U_{11} & U_{12} \\ U_{21} & U_{22} \end{pmatrix} \cdot \begin{pmatrix} I & 0 \\ 0 & T_2 \end{pmatrix} \;=\; \begin{pmatrix} V_{11} & 0 \\ 0 & V_{22} \end{pmatrix}.$$

Here $V_{11} = \mathrm{diag}(v_{11}, \ldots, v_{rr})$ agrees with the upper $r \times r$-block of V, and V_{22} is a diagonal matrix whose entries are bounded above by u in absolute value. This implies $U_{11} = V_{11}$, $U_{12} = 0$, $U_{21} = 0$, and U_{22} is an arbitrary $(m-r)\times(n-r)$ matrix with spectral norm at most u. Hence the Voronoi cell of V is congruent to the set of all such matrices U_{22}. This convex body equals u times the unit ball in $\mathbb{R}^{(m-r)\times(n-r)}$ under the spectral norm. $\square$

Our problem becomes even more interesting when we restrict to matrices in a linear subspace. To see this, let X denote the variety of symmetric $n \times n$ matrices of rank $\leq r$. We can regard X either as a variety in the ambient matrix space $\mathbb{R}^{n\times n}$, or in the space $\mathbb{R}^{\binom{n+1}{2}}$ whose coordinates are the upper triangular entries of a symmetric matrix. On the latter space we have both the *Euclidean norm* and the *Frobenius norm*. These are now different!

The Frobenius norm on $\mathbb{R}^{\binom{n+1}{2}}$ is the restriction of the Frobenius norm on $\mathbb{R}^{n\times n}$ to the subspace of symmetric matrices. For instance, if $n = 2$, we identify the vector (a, b, c) with the symmetric matrix $\left(\begin{smallmatrix} a & b \\ b & c \end{smallmatrix}\right)$. The Frobenius norm is $\sqrt{a^2+2b^2+c^2}$, whereas the Euclidean norm is $\sqrt{a^2+b^2+c^2}$. The two norms have dramatically different properties with respect to low rank approximation. The Eckart-Young Theorem remains valid for the Frobenius norm on $\mathbb{R}^{\binom{n+1}{2}}$, but this is not valid for the Euclidean norm (cf. [8, Example 3.2]). In what follows we elucidate this point by comparing the Voronoi cells with respect to the two norms.

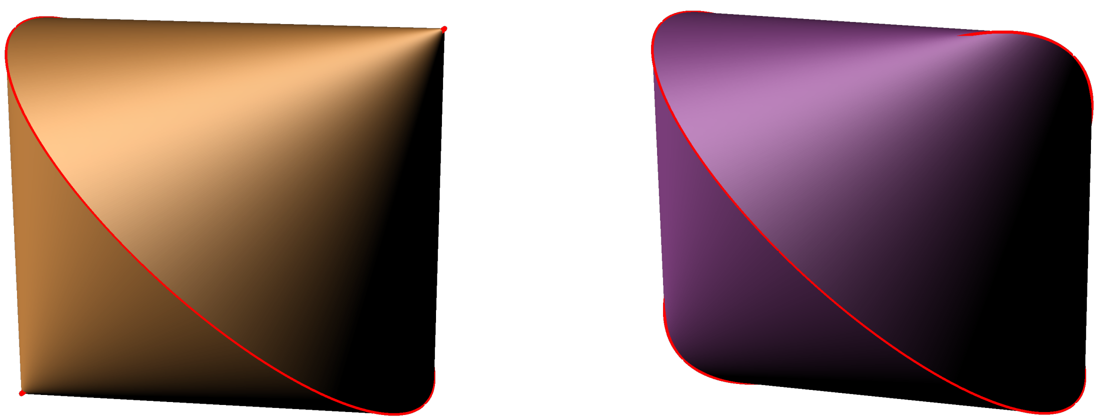

Figure 3: The Voronoi cell of a symmetric 3×3 matrix of rank 1 is a convex body of dimension 3. It is shown for the Frobenius norm (left) and for the Euclidean norm (right).

Example 22. Let X be the variety of symmetric 3×3 matrices of rank ≤ 1. For the Euclidean metric, X lives in $\mathbb{R}^6$. For the Frobenius metric, X lives in a 6-dimensional subspace of $\mathbb{R}^{3\times 3}$. Let V be a regular point in X, i.e. a symmetric 3×3 matrix of rank 1. The normal space to X at V has dimension 3. Hence, in either norm, the Voronoi cell $\mathrm{Vor}_X(V)$ is a 3-dimensional convex body. Figure 3 illustrates these bodies for our two metrics.

For the Frobenius metric, the Voronoi cell is congruent to the set of matrices $\left(\begin{smallmatrix} a & b \\ b & c \end{smallmatrix}\right)$ with eigenvalues between -1 and 1. This semialgebraic set is bounded by the surfaces defined by the singular quadrics $\det\left(\begin{smallmatrix} a+1 & b \\ b & c+1 \end{smallmatrix}\right)$ and $\det\left(\begin{smallmatrix} a-1 & b \\ b & c-1 \end{smallmatrix}\right)$. The Voronoi ideal is of degree 4, defined by the product of these two determinants (modulo the normal space). The Voronoi cell is shown on the left in Figure 3. It is the intersection of two quadratic cones. The cell is the convex hull of the circle in which the two quadrics meet, together with the two vertices.

For the Euclidean metric, the Voronoi boundary at a generic point V in X is defined by an irreducible polynomial of degree 18 in a, b, c. In some cases, the Voronoi degree can drop. For instance, consider the special rank 1 matrix $V = \left(\begin{smallmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{smallmatrix}\right)$. For this point, the degree of the Voronoi boundary is only 12. This particular Voronoi cell is shown on the right in Figure 3. This cell is the convex hull of two ellipses, which are shown in red in the diagram.

References

[1] Y. Alexandr and A. Heaton: *Logarithmic Voronoi cells*, Algebraic Statistics **12** (2021) 75–95.

[2] A. Becedas, K. Kohn and L. Venturello: *Voronoi diagrams of algebraic varieties under polyhedral norms*, `arXiv:2209.11463`.

[3] M. Brandt and M. Weinstein: *Voronoi cells in metric algebraic geometry of plane curves*, `arXiv:1906.11337`.

[4] P. Breiding, S. Kalisnik, B. Sturmfels and M. Weinstein: *Learning algebraic varieties from samples*, Revista Matematica Complutense **31** (2018) 545–593.

[5] D. Cifuentes, C. Harris and B. Sturmfels: *The geometry of SDP-exactness in quadratic optimization*, Mathematical Programming, Ser. A **182** (2020) 399–428.

[6] D. Cifuentes, K. Ranestad, B. Sturmfels and M. Weinstein: *Voronoi cells of varieties*, Journal of Symbolic Computation **109** (2022) 351–366.

[7] S. Di Rocco, D. Eklund and M. Weinstein: *The bottleneck degree of algebraic varieties*, SIAM J. Appl. Algebra Geom. **4** (2020) 227–253.

[8] J. Draisma, E. Horobeţ, G. Ottaviani, B. Sturmfels and R. Thomas: *The Euclidean distance degree of an algebraic variety*, Found. Comput. Math. **16** (2016) 99–149.

[9] E. Dufresne, P. Edwards, H. Harrington and J. Hauenstein: *Sampling real algebraic varieties for topological data analysis*, in 18th IEEE International Conference on Machine Learning and Applications (ICMLA 2019).

[10] D. Eklund: *The numerical algebraic geometry of bottlenecks*, Advances in Applied Mathematics **142** (2023), Paper No. 102416, 20pp.

[11] W. Fulton: *Intersection Theory*, Springer Verlag, 1984.

[12] D. Grayson and M. Stillman: *Macaulay2, a software system for research in algebraic geometry*, available at `www.math.uiuc.edu/Macaulay2/`.

[13] E. Horobeţ and M. Weinstein: *Offset hypersurfaces and persistent homology of algebraic varieties*, Comput. Aided Geom. Design **74** (2019), 101767, 14 pp.

[14] G. Ottaviani and L. Sodomaco: *The distance function from a real algebraic variety*, Comput. Aided Geom. Design **82** (2020) 101927, 20 pp.