K-transversals of parallel convex sets

Nina Amenta
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304, USA
amenta@parc.xerox.com

May 6, 1996

Abstract

$\mathbb{R}^d$ can be divided into a union of *parallel* $(d-k)$-flats of the form $x_1 = g_1, x_2 = g_2, \ldots x_k = g_k$, where the g_i are constant. Let C be a family of parallel $(d-k)$-dimensional convex sets, meaning that each is contained in one of the above parallel $(d-k)$-flats. We give a parameterization of the set of k-flats in $\mathbb{R}^d$, such that the set of k-flats which intersect, in a point, any set $c \in C$, is convex. Parameterizing the lines in $\mathbb{R}^3$ through horizontal convex sets as convex sets has applications to medical imaging, and interesting connections with recent work on light field rendering in computer graphics. The general case is useful for fitting k-flats to points in $\mathbb{R}^d$.

The following easy reduction is well known. Let C be a finite set of parallel line segments in $\mathbb{R}^d$. We want to find a $(d-1)$-*transversal* for C, that is, a hyperplane intersecting every segment in C. Such a hyperplane has to pass below the upper endpoint of each segment and above the lower endpoint. In the dual, the endpoints correspond to linear halfspaces, and the intersection of these halfspaces corresponds to the set of hyperplane transversals of the parallel segments in the primal. So the problem is solved by linear programming in dimension d, in linear time if d is fixed.

Here, we give the appropriate generalization of this observation for k-transversals, for $1 \leq k < d$. A k-transversal of a family of sets C is a k-flat (that is, a k-dimensional affine subspace) intersecting every set in C. Figure 1 shows the case $k = 1, d = 3$.

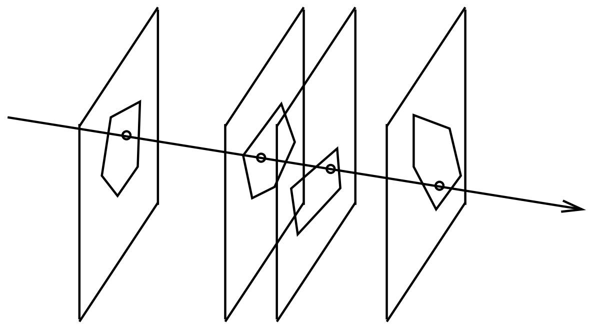

Figure 1: The set of lines intersecting all members of a family of parallel polygons can be represented as a convex set.

A family C of $(d-k)$-dimensional sets in $\mathbb{R}^d$ are *parallel* if they can be rotated so that each set $c \in C$ lies in a $(d-k)$-flat

$$x_1 = g_1, x_2 = g_2, \ldots x_k = g_k$$

where $g_1, \ldots g_k$ are constants and $x_1, \ldots, x_k$ are the first k coordinates of a point $x \in \mathbb{R}^d$. From now on we will just assume that C is so rotated. We say that a k-flat y intersects a set $c \in C$ *non-degenerately* if $y \cap c$ consists of a single point; a k-flat in general position intersects a set $c \in C$ non-degenerately, if at all. Our Main Theorem 2.2 gives a parameterization under which the k-flats intersecting non-degenerately any member of C form a convex set in $\mathbb{R}^{(k+1)(d-k)}$. This result is a simple algebraic consequence of adopting the "right" parameterization of k-flats in $\mathbb{R}^d$. But it has both mathematical and practical implications.

1 Background

An immediate consequence is the following.

Theorem 1.1 *The Helly number for k-transversals of parallel $(d-k)$-dimensional convex sets is $(k+1)(d-k)+1$.*

A family of sets $\mathbb{C}$ has *Helly number h* for some property Π (here, the property of having a k-transversal) when h is the smallest integer (if one exists) such that any finite subfamily $C \subseteq \mathbb{C}$ has property Π if and only if every subfamily $B \subseteq C$ with $|B| \leq h$ also has Π. Theorems of the form "$\mathbb{C}$ has Helly number h" are called *Helly-type theorems* because they follow the model of Helly's theorem, which states that the family of convex sets in R^d has Helly number $d+1$. Helly's theorem together with Theorem 2.2 implies Theorem 1.1.

There are many Helly-type theorems about hyperplane transversals, and some about line transversals (see [GPW93]), but this is the first theorem giving a finite Helly number for k-transversals for all $1 \leq k < d$ for some family of sets. While, as we observed in the introduction, the space of hyperplanes in $\mathbb{R}^d$ is isomorphic to $\mathbb{R}^d$, the space of k-flats in $\mathbb{R}^d$, for $1 \leq k < (d-1)$, is a curved projective manifold, known as a *Grassmanian*, and generally much more difficult to work with. Goodman and Pollack define the convex sets in a Grassmanian as the sets of k-transversals of convex sets in $\mathbb{R}^d$ [GP94]. They show that this definition is a generalization of convexity, in some senses, but, for instance, convex set under their definition can have multiple connected components. We exhibit subsets of the Grassmanian which are convex in the usual sense.

The special case of Theorem 1.1 for $k = 1, d = 3$, was given by Grünbaum [G60], which suggested our approach.

One immediate algorithmic consequence of Theorem 2.2 is that a k-transversal of a finite set C of parallel $(d-k)$-dimensional sets, if one exists, can be found by a convex program in dimension $(k+1)(d-k)$, in linear expected time if d is constant. The case of line transversals in $\mathbb{R}^3$ is the first interesting one, and it has some applications in computer graphics.

Medical images, such as CAT scans and MRI images, are given as a set of parallel two-flats. When the regions in each image are decomposed into convex pieces, the line transversals of the various possible subsets of pieces form a convex subdivision of $\mathbb{R}^4$ under our parameterization. This subdivision is interesting for volume visualization and, as we discuss below, for path planning for lasers, needles or other invasive linear elements.

Our parameterization is also used in computer graphics in a recent paper by Hanrahan and Levoy on *light field rendering* [HL96], in which an object is represented by the radiance on the directed lines incident to it. A light field is a hyper-rectangular set of lines including those incident to the object. A quantized light field is stored in a four-dimensional array, and an image is constructed by selecting the lines through a particular viewpoint, which correspond to a two-flat through the hyper-rectangle. Theorem 2.1 implies that certain linear halfspaces through the hyper-rectangle correspond to the sets of lines of constant depth in $\mathbb{R}^3$. This may have some application to the important problem of reconstructing the three-dimensional shape of an object from its light field representation.

Finding line transversals is an important subproblem in the visibility preprocessing of large scenes in computer graphics, although admittedly it is difficult to imagine a scene which requires solving the special case of the problem treated here.

In general dimension, Theorem 2.2 can be applied to the problem of fitting k-flats to points. Every point $x \in \mathbb{R}^d$ is contained in exactly one member g of a set of parallel $(d-k)$-flats, and any k-flat f in general position intersects g in exactly one point. We define the $(d-k)$-*dimensional distance* between x and f to be the Euclidean distance from x to the point in which f intersects g. Figure 2 again shows the case $k = 1, d = 3$. This metric is a higher-dimensional analogue of

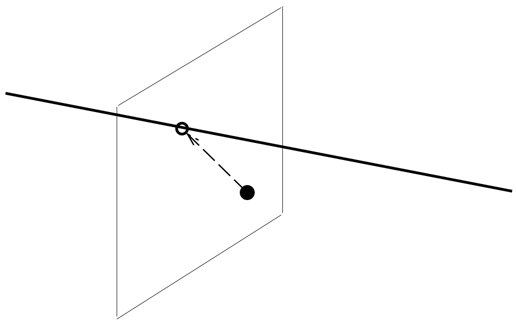

Figure 2: The distance between f and x is measured in the $(d-k)$-flat g.

the vertical distance. We show that finding the k-flat which minimizes the maximum $(d-k)$-dimensional distance to any member of a set of points is a Convex Programming problem and can be solved in expected linear time.

2 Main Theorem

We have defined *parallel* $(d-k)$-flats in $\mathbb{R}^d$ to be the affine subspaces satisfying $x_1 = g_1, x_2 = g_2, \ldots x_k = g_k$, where the g_i are constants. Let us define the set $\mathbb{H}$ of parallel $(d-k)$-*half-flats* to be the sets formed by the intersection of a linear halfspace in $\mathbb{R}^d$ with one of the parallel $(d-k)$-flats, that is, a set satisfying the equalities

$$\begin{array}{rcl} x_1 & = & g_1 \\ & \cdots & \\ x_k & = & g_k \end{array}$$

and some inequality

$$a_1 x_{k+1} + \ldots a_{d-k} x_d \geq 1$$

A $(d-k)$-half-flat $h_{a,g}$ in $\mathbb{H}$ is determined by the g_i and the a_j, and so has d coefficients.

A k-flat in general position in $\mathbb{R}^d$ intersects any $(d-k)$-flat g in a single point. We will parameterize a k-flat in general position by its points of intersection with each of the $k+1$

following $(d-k)$-flats:

$$\begin{aligned}
u_0 &= \{x_1 = 0, x_2 = 0, \ldots, x_k = 0\} \\
u_1 &= \{x_1 = 1, x_2 = 0, \ldots, x_k = 0\} \\
u_2 &= \{x_1 = 0, x_2 = 1, \ldots, x_k = 0\} \\
&\ldots \\
u_k &= \{x_1 = 0, x_2 = 0, \ldots, x_k = 1\}
\end{aligned}$$

Such a point, for each u_0, is specified by the $(d-k)$ values of $x_{k+1}, \ldots, x_d$, which we shall call $y_{(i,k+1)}, \ldots, y_{(i,d)}$. A k-flat Y in general position is thus specified by $(k+1)(d-k)$ independent parameters, the entries in the matrix

$$Y = \begin{bmatrix} y_{(0,k+1)} & \cdots & y_{(0,d)} \\ & \cdots & \\ y_{(k,k+1)} & \cdots & y_{(k,d)} \end{bmatrix}$$

In order to simplify the notation below, we will subtract the first row from each of the subsequent rows, representing a k-flat by the matrix

$$Y' = \begin{bmatrix} y_{(0,k+1)} & \cdots & y_{(0,d)} \\ y_{(1,k+1)} - y_{(0,k+1)} & \cdots & y_{(1,d)} - y_{(0,d)} \\ & \cdots & \\ y_{(k,k+1)} - y_{(0,k+1)} & \cdots & y_{(k,d)} - y_{(0,d)} \end{bmatrix}$$

This corresponds to an affine transformation of the space of k-flats. Rows $1, \ldots k$ now express the change in $x_{k+1}, \ldots, x_d$ per unit change in $x_1, \ldots, x_k$.

Theorem 2.1 *The $(k+1)(d-k)$-dimensional set of k-flats in $\mathbb{R}^d$ can be parameterized so that the k-flats intersecting, non-degenerately, any $(d-k)$-dimensional half-flat in $\mathbb{H}$ form a linear halfspace in $\mathbb{R}^{(k+1)(d-k)}$.*

Proof. The points of the k-flat Y are the points of the form:

$$\begin{pmatrix} x_1, \ldots, x_k \\ (1, x_1, \ldots x_k) \cdot Y' \end{pmatrix}$$

This notation indicates the concatenation of $x_1, \ldots, x_k$ with the vector $(1, x_1, \ldots x_k) \cdot Y'$. A k-flat in general position will intersect the $(d-k)$-flat $x_1 = g_1, x_2 = g_2, \ldots, x_k = g_k$ in the point

$$\begin{pmatrix} g_1, \ldots, g_k \\ (1, g_1, \ldots g_k) \cdot Y' \end{pmatrix}$$

That point lies in the half-flat $h_{a,g}$ if and only if

$$(1, g_1, \ldots, g_k) \cdot Y' \cdot (a_1, \ldots, a_{(d-k)})^T \geq 1$$

Since the g_i and the a_j are constants, $h_{a,g}$ induces a linear inequality on the $y'_{(i,j)}$, and therefore also on the $y_{(i,j)}$. $\square$

Theorem 2.2 *The $(k+1)(d-k)$-dimensional set of k-flats in $\mathbb{R}^d$ can be parameterized so that the non-degenerate k-transversals of any family C of parallel $(d-k)$-dimensional convex sets form a convex set in $\mathbb{R}^{(k+1)(d-k)}$.*

Proof. A convex set $c \in C$ is the intersection of a (possibly infinite) family H of $(d-k)$-dimensional half-flats in $\mathbb{H}$. By Theorem 2.1, the set of k-flats non-degenerately intersecting such a half-flat form a linear halfspace in $\mathbb{R}^{(k+1)(d-k)}$. The k-flats intersecting *every* half-flat in H correspond to the intersection c' of the corresponding halfspaces in $\mathbb{R}^{(k+1)(d-k)}$. This is a convex set. So the intersection of the c' is an intersection of convex sets, and so again a convex set. □

3 Some algorithmic corollaries

We describe a few of the algorithmic implications of our main theorems.

3.1 Finding k-transversals

From Theorem 2.1, we can infer immediately

Corollary 3.1 *A non-degenerate k-transversal of a family C of parallel $(d-k)$-dimensional polytopes, if one exists, can be found by linear programming in dimension $(k+1)(d-k)$. When d is fixed, this requires time linear in the total number of facets of C.*

And, from Theorem 2.2,

Corollary 3.2 *Let C be a finite family of parallel $(d-k)$-dimensional convex sets. A non-degenerate k-transversal of C, if one exists, can be found by convex programming in dimension $(k+1)(d-k)$.*

Convex programming is the problem of minimizing a convex objective function over the intersection of a family of convex sets. Any convex function on the space of k-flats can be used as the objective function for the convex program in Theorem 3.2, most conveniently a linear function. Convex programming is an *LP-type problem*, as defined in [MSW92] (see [A94b] for a little more on convex programming). This means that if d is constant and the k-flat minimizing the objective function for any subset of at most $(k+1)(d-k)+1$ members of C, if one exists, can be found in time t_b, then a line transversal for C can be found in expected time $O(|C| + t_b \lg |C|)$, which is linear in $|C|$ when t_b is small enough.

3.2 Medical images and path planning

Medical images of three-dimensional anatomy such as CAT scans and MRI images are given as intensity images in a family of parallel slices. These slices can automatically segmented so that each is represented as a union of polygonal convex regions of constant or continuously varying intensity. Each region is assumed to represent a slice of a particular kind of tissue. Under our parameterization, the lines bounding these polygons correspond to an arrangement of hyperplanes in the four-dimensional space of lines in $\mathbb{R}^3$.

We sketch one algorithmic consequence of this observation. Consider the problem of finding a path for a biopsy needle which goes to a tumor while missing a collection of vital organs. The vital organs, the non-vital tissue, and the tumor are all represented by collections of parallel convex polygons. We wish to find all acceptable paths for the needle. The set of acceptable paths corresponds to a union of cells in the corresponding hyperplane arrangement in the four-dimensional space of lines. Each cell in this arrangement is a subset of lines.

The arrangement can be constructed by random sampling [C87]. We select a constant size random sample of the parallel polygons, construct the arrangement induced by their edges in the space of lines, and subdivide each cell of this arrangement into simplices. We construct a subproblem for each simplex consisting of the polygons which intersect any of the lines in the simplex. For each simplex s, we maintain the set of polygons for which are intersected by every line in s. These can be kept in sorted order, since the polygons are parallel. We also maintain the first polygon in this set corresponding to a vital organ, if any, and the first corresponding to the tumor. Recursively proceeding on the subproblems gives us a tree which represents the arrangement. A leaf in this tree is a set of lines, and tracing the path from the leaf to the root gives us all the polygons intersected by that set of lines.

This data structure requires time and space $O(n^{4+\varepsilon})$. To find the leaf cells corresponding to acceptable needle paths, we traverse the arrangement by depth-first search and keep track of whether the tumor or a vital organ is hit first by the current set of lines. If a path exists, we will find at least one leaf for which this is true. While this gives an $O(n^{4+\varepsilon})$ algorithm, our intuition is that in practice it would be much more efficient. The set of lines passing through three polygon edges is unlikely to intersect a fourth edge, so in practice the algorithm as described will probably run in roughly $O(n^3)$ time. Furthermore, the only important cells are those intersecting the tumor. Only constructing those cells could reduce the running time to something like $O(n^2)$. The fact that the representation is linear makes the algorithm feasible to implement.

3.3 Fitting k-flats to points

We defined the $(d-k)$-dimensional distance from a point x to a k-flat f in general position to be the Euclidean distance from x to the point in which f intersects the unique member g of the set of parallel $(d-k)$-flats containing x. This metric is not as exotic as it may seem. When we fit a k-flat to a set of points using Least Squares, we are computing the k-flat which minimizes the sum of the squared $(d-k)$-dimensional distances to each of the points. The metric is appropriate, for instance, when x is a multidimensional data point for which the coordinates $x_1, \ldots, x_k$ represent variables which are known exactly and $x_{k+1}, \ldots, x_d$ represent variables which are measured with some error. Here, we use combinatorial methods to compute the k-flat which *minimizes the maximum* $(d-k)$-dimensional distance to any point, in time linear in the number of points.

Let X be our set of n points in $\mathbb{R}^d$. The region at $(d-k)$-dimensional distance at most ε from a point $x \in X$ is a $(d-k)$-dimensional disk c_ε in the unique flat g containing x in the set of parallel $(d-k)$-flats. For the entire set X, these disks form a set $C_\varepsilon = \{c_\varepsilon \parallel x \in X\}$ of parallel $(d-k)$-dimensional convex polytopes. Now consider the $(k+1)(d-k)+1$ dimensional cross-product $Y \times \mathbb{R}^+$. A point y, ε in this space represents a k-flat in $\mathbb{R}^d$ and a value of ε.

Lemma 3.3 *$Y \times \mathbb{R}^+$ can be parameterized so that, for any point $x \in X$, the set of points y, ε which correspond to a k-flat intersecting, non-degenerately, the disk c_ε around x form a convex set.*

Proof. Each disk c_ε is the intersection in g of an infinite family $H_{g,\varepsilon} \subset \mathbb{H}$ of $(d-k)$-dimensional half-flats, each $h_{a,g,\varepsilon}$ of the form

$$a_1 x_{k+1} + \ldots a_{d-k} x_d - \varepsilon \geq a_{d-k+1}$$

where the a_i are normalized so that $a_1^2 + a_2^2 + \ldots a_{d-k}^2 = 1$, and a_{d-k+1} is determined by the requirement that at $\varepsilon = 0$, the equality $a_1 x_{k+1} + \ldots a_{d-k} x_d = a_{d-k+1}$ will be satisfied by the point $x \in X$ at the center of the disk.

Any k-flat Y' in general position intersects g in the point $g \cdot Y'$, and that point lies in $h_{a,g,\varepsilon}$ if and only if

$$(1, g_1, \ldots, g_k) \cdot Y' \cdot (a_1, \ldots, a_{(d-k)})^T - \varepsilon \geq a_{d-k+1}$$

This is a linear inequality in the $Y' \times \mathbb{R}^+$. The set of k-flats intersecting every c_ε is the intersection of this infinite set $H_{g,\varepsilon}$ of linear halfspaces, and hence a convex set. □

So the sets of close-enough k-flats at every ε form convex sets of points in $R^{(k+1)(d-k)+1}$. To find the minimum ε at which there is a k-flat that is within ε of every point, we just have to minimize the linear function ε over the intersection of these convex sets. This gives us the following.

Theorem 3.4 *Let X be a finite family of points in $\mathbb{R}^d$. The k-flat which minimizes the maximum $(d-k)$-dimensional distance to any point of X can be found by convex programming in dimension $(k+1)(d-k)+1$, in linear time when d is fixed.*

Note that this result also applies to distance functions in which the disk around every point is replaced by some other $(d-k)$-dimensional unit ball, for example what we might call the $(d-k)$-dimensional L^1 distance or the $(d-k)$-dimensional L^∞ distance.

4 Remarks

4.1 Disclaimer

Note that these theorems only apply to non-degenerate k-transversals. If the parallel $(d-k)$-sets in C fail to span $\mathbb{R}^d$, they might have a degenerate k-transversal, which intersects some $c \in C$ in a subspace of dimension greater than zero. In the first interesting case of line transversals in $\mathbb{R}^3$, there may be a degenerate transversal when the parallel two-dimensional convex sets in C all lie in the same plane. Finding a line transversal of a family of convex sets in the plane is clearly not a convex programming problem, since the set of line transversals may have up to n connected components, $n = |C|$. And in fact there is a lower bound of $\Omega(n \lg n)$ for the special case of finding a line transversal for a family of unit balls in the plane [LW86].

4.2 Projective transformation

The family of parallel $(d-k)$-flats can be defined as the set of $(d-k)$-flats intersecting a $(d-k-1)$-flat at infinity f_∞ spanned by the points at infinity on the $k+1, \ldots, d$ axes. Theorem 2.1 tells us that the family of k-flats intersecting a $(d-k-1)$-flat contained in one of these parallel $(d-k)$-flats forms a hyperplane under our parameterization. These $(d-k-1)$-flats also intersect f_∞. Consider a projective transformation which moves f_∞ to an arbitrary $(d-k-1)$-dimensional flat f.

Corollary 4.1 *Let H_f be the set of $(d-k-1)$-flats intersecting a given $(d-k-1)$-flat f in $\mathbb{R}^d$. The set of k-flats in $\mathbb{R}^d$ can be parameterized so that the k-flats intersecting any $(d-k-1)$-dimensional flat in H_f form a hyperplane in $\mathbb{R}^{(k+1)(d-k)}$.*

4.3 Axis aligned boxes

The following easy observation is a special case of Theorem 2.2.

Observation 4.2 *The k-flats in $\mathbb{R}^d$ can be parameterized so that the set of k-flats intersecting any member of a family of parallel $(d-k)$-dimensional axis-aligned boxes is convex.*

Such an axis-aligned box can be defined as points satisfying

$$\begin{aligned} x_1 &= g_1 \\ &\cdots \\ x_k &= g_k \end{aligned}$$

and the inequalities

$$\begin{aligned} &a_1 \leq x_{k+1} \leq b_1 \\ &\cdots \\ &a_{(d-k)} \leq x_d \leq b_{(d-k)} \end{aligned}$$

Substituting in the expression for the intersection of Y' and g, we get

$$(a_1, \ldots, a_{(d-k)})^T \ \leq \ (1, g_1, \ldots, g_k) \ \cdot \ Y' \ \leq \ (b_1, \ldots, b_{(d-k)})^T$$

This system can be separated into $(d-k)$ separate systems of linear inequalities, one for each column of Y', and solved as $(d-k)$ lower-dimensional linear programs, which is much faster.

References

[A94b] Nina Amenta, Helly Theorems and Generalized Linear Programming, *Discrete and Computational Geometry* 12:3 (1994) pages 241-261.

[C87] Kenneth L. Clarkson. New Applications of random sampling in computational geometry, *Discrete and Computational Geometry* **2** (1987) pp 195-222

[GPW93] Jacob E. Goodman, Richard Pollack and Rephael Wenger. Geometric Transversal Theory, in New Trends in Discrete and Computational Geometry, János Pach, ed., (1993) Springer Verlag, Berlin, pages 163-198.

[GP94] Jacob E. Goodman and Richard Pollack. Foundations of a theory of convexity on affine grassmann manifolds, *manuscript*, (1994)

[G60] Branko Grünbaum. Common transversals for families of sets, *Journal of the London Mathematical Society* **35** (1960), pages 408-16.

[HL96] Pat Hanrahan and Mark Levoy. Light field rendering, *Siggraph 96*, to appear.

[LW86] D.T. Lee and Y. F. Wu. Geometric complexity of some location problems, *Algorithmica* **1** (1986), pages 193-211.

[MSW92] Jiří Matoušek, Micha Sharir and Emo Welzl. A subexponential bound for linear programming, *Proceedings of the 8th Annual Symposium on Computational Geometry* (1992) pages 1-8.