Smoothed Analysis of Probabilistic Roadmaps

Siddhartha Chaudhuri[*,1], Vladlen Koltun[2]

Computer Science Department, Stanford University, 353 Serra Mall, Stanford, CA 94305, USA

Abstract

The probabilistic roadmap algorithm is a leading heuristic for robot motion planning. It is extremely efficient in practice, yet its worst case convergence time is *unbounded* as a function of the input's combinatorial complexity. We prove a smoothed polynomial upper bound on the number of samples required to produce an accurate probabilistic roadmap, and thus on the running time of the algorithm, in an environment of simplices. This sheds light on its widespread empirical success.

Key words: motion planning, probabilistic roadmap, smoothed analysis, locally orthogonal decomposition

1. Introduction

Smoothed analysis. It is well-documented that many geometric algorithms that are extremely efficient in practice have exceedingly poor worst-case performance guarantees. Two approaches were put forth to address this issue. The first tries to formally model various classes of inputs that arise in practice and analyze the performance of algorithms on these models [16]. For example, it was proposed that in practice geometric objects are *fat* [1, 13, 32, 40], point sets have bounded *spread* [8, 10, 18, 19], and geometric scenes have *low density*, are *uncluttered*, *sparse*, etc. [6, 14, 15, 34].

The second approach stems from the observation that geometric inputs often contain a small amount of random noise, such as with point clouds generated by a laser scanner [30]. It can be argued that small degrees of randomness creep into geometric inputs even if they are created by a human modeler [37]. By this reasoning, finely tuned worst-case examples have a low probability of arising and should not disproportionately skew theoretical measures of algorithm performance. This is formalized in smoothed analysis [39], which measures the maximum over inputs of the expected running time of the algorithm under slight random perturbations of those inputs. For example, let $A \in \mathbb{R}^{n\times d}$ specify a set of n points in $\mathbb{R}^d$, and let $f_X(A)$, where $f_X : \mathbb{R}^{n\times d} \mapsto \mathbb{R}$, be a measure of the performance of algorithm X on A. Then the worst-case performance of X is

$$\max_{A\in\mathbb{R}^{n\times d}} f_X(A),$$

[*]Corresponding author.

Email address: `sidch@cs.stanford.edu` (Siddhartha Chaudhuri)

[1]Supported by a PACCAR Inc. Stanford Graduate Fellowship.

[2]Supported by NSF CAREER award CCF-0641402 and an Alfred P. Sloan Research Fellowship.

the average-case performance of X is

$$\mathop{\mathrm{E}}_{A \sim \mathcal{D}} \left[f_X(A) \right],$$

where $\mathcal{D} : \mathbb{R}^{n \times d} \mapsto \mathbb{R}$ is a suitable distribution, and the smoothed performance of X is

$$\max_{A \in \mathbb{R}^{n \times d}} \mathop{\mathrm{E}}_{R \sim \mathcal{N}} \left[f_X(A + \|A\| R) \right],$$

where $\|A\|$ denotes the Frobenius norm of A or some similar measure of the *numerical* magnitude of the input, and $\mathcal{N} = N(0, \sigma^2 I_{n \times d})$ is a Gaussian distribution in $\mathbb{R}^{n \times d}$ with mean 0 and variance σ^2. The parameter σ controls the magnitude of the random perturbation, and as it varies from 0 to ∞ the smoothed performance measure interpolates between worst-case and average-case performance.

Smoothed analysis is a new framework that has already been applied to a wide variety of problems [3, 4, 7, 11, 12, 17, 38]. Its advantage compared to the above-described explicit formulation of realistic input models lies in its generality and immediate applicability across contexts, and its reliance on only one assumption, namely the presence of some degree of randomness in the input.

Probabilistic roadmaps. The probabilistic roadmap (PRM) algorithm revolutionized robot motion planning [23, 25, 27]. It is a simple heuristic that exhibits rapid performance and has become the standard algorithm in the field [20, 21, 36]. Yet its worst-case running time is *unbounded* as a function of the input's combinatorial complexity. The basic algorithm for constructing a probabilistic roadmap is as follows:

> Sample uniformly at random a set of points, called milestones, from the *configuration space* $\mathcal{C}$ of the robot. Keep only those milestones that lie in the *free configuration space* $\mathcal{C}_{\mathsf{free}}$. [3] Let V be the resulting point set. For every $u, v \in V$, if the straight line segment between u and v lies entirely in $\mathcal{C}_{\mathsf{free}}$, add $\{u, v\}$ to the set of edges E, initially empty. The graph $G = (V, E)$ is the probabilistic roadmap.

Given such a roadmap G, a motion between two points p, q in $\mathcal{C}_{\text{free}}$ can be constructed as follows:

> Find a milestone p' (resp., q') in V that is visible from p (resp., from q). If p' and q' lie in different connected components of G, report that there is no feasible motion between p and q. Otherwise plan the motion using a path in G that connects p' and q'.

The above PRM construction and query algorithms can be efficiently implemented in very general settings. The outstanding issue is what the number of samples should be to guarantee (in expectation) that G accurately represents the connectivity of $\mathcal{C}_{\text{free}}$. Clearly,

[3] A robot's *configuration space* is the set of physical positions it may attain (which may or may not coincide with obstacles), parametrized by its degrees of freedom (so a robot with d degrees of freedom has a d-dimensional configuration space). The robot's *free configuration space* is the subset of these positions which do *not* coincide with obstacles, i.e. are possible in real life. These terms are standard in the motion planning literature [29].

for the algorithm to be accurate there should be a milestone visible from any point in $\mathcal{C}_{\text{free}}$, and there should be a bijective correspondence between the set of connected components of G and the set of connected components of $\mathcal{C}_{\text{free}}$. Unfortunately, the number of random samples required to guarantee this can be made arbitrarily large even for very simple configuration spaces [21].

A number of theoretical analyses provide bounds for the number of samples under assumptions on the structure of $\mathcal{C}_{\text{free}}$ such as goodness [5, 26], expansiveness [22], and the existence of high-clearance paths [24]. However, none of these assumptions were justified in terms of realistic motion planning problems. In practice, the number of random samples is chosen ad hoc.

Contributions. This paper initiates the use of smoothed analysis to explain the success of PRM. We model the free configuration space of the robot using a set of n $(d-1)$-simplices in $\mathbb{R}^d$, which act as obstacles. The vertices of these simplices are subject to Gaussian perturbations of variance $(\sigma D)^2$, where D is the diameter of the configuration space. We prove a smoothed upper bound on the required number of milestones that is polynomial in n and $\frac{1}{\sigma}$. The result extends to all γ-smooth perturbations, see below.

In order to achieve this bound we define a space decomposition called the locally orthogonal decomposition. Previously known decompositions, like the vertical decomposition [9, 28] and the "castles in the air" decomposition [2] turn out to be unsuitable for our purpose. We prove that for the roadmap to accurately represent the free configuration space it is sufficient that a milestone is sampled from every cell of this decomposition. We then prove a smoothed lower bound on the volume of every decomposition cell. This leads to the desired bound on the number of milestones.

Our result is only a step towards a convincing theoretical justification of PRM. The analysis is quite challenging already for the simple representation of the configuration space using independently perturbed simplices. In Section 4 we outline directions for its extension to more general configuration space models.

2. Bounding the Number of Milestones

Notation. Let V be a d-dimensional vector space and assume d to be constant. For $0 \leq k \leq d$, a *k-subspace* of V is the set of linear combinations of k linearly independent vectors lying in V. A subspace necessarily contains the origin. A *k-flat* is an affine translation of a k-subspace. Points are 0-flats, straight lines are 1-flats, planes are 2-flats and hyperplanes are $(d-1)$-flats. A k-dimensional flat is the intersection of $d-k$ hyperplanes.

A hyperplane divides V into two *halfspaces*. More generally, a set of hyperplanes $\mathcal{H}$ subdivides V into a number of disjoint, open, d-dimensional *cells*. Further, assume a subset $\mathcal{U}$ of $\mathcal{H}$ intersects in a k-flat F, and let $\mathcal{U}'$ be the set of hyperplanes in $\mathcal{H}$ which intersect F but do not contain it. $\mathcal{U}'$ subdivides F into disjoint, open, k-dimensional regions called k-faces (if $\mathcal{U}'$ is empty, F is a k-face on its own — note that this handles the special case of 0-faces, which are closed, not open). A 0-face is called a *vertex*, and a $(d-1)$-face is called a *facet*. Extending the notation, a cell is considered a d-face. The entire structure is referred to as the *arrangement* of the set of hyperplanes $\mathcal{H}$. An arrangement of hyperplanes is a *convex subdivision*, since all its faces are convex sets.

A set of hyperplanes $\mathcal{H}$ (or their arrangement) is in *general position* if every d-tuple of hyperplanes in $\mathcal{H}$ intersect in exactly one point. We note that the precise meaning of "general position" we adopt here defines a suitable "general case" for our problem — other authors may use different notions.

The *distance* between two flats X and Y is defined as $\min_{x \in X, y \in Y} \|x - y\|$. Two flats are said to be ε*-close* if their distance is at most ε; otherwise they are ε*-distant*.

$A \oplus B$ is the *Minkowski sum* of sets A and B, i.e. it is the set of all sums of the form $a + b$, where $a \in A$ and $b \in B$.

The d*-ball* of radius r, denoted $B_d(r)$, is the set of points at distance at most r from the origin. ($B_d(x, r)$ is defined as $B_d(r) \oplus x$.) The boundary of $B_d(r)$ is the $(d-1)$*-sphere* of radius r, written $S_{d-1}(r)$.

The *volume* of a k-dimensional object will refer to its k-dimensional Lebesgue measure. If this object is embedded in a space of higher dimension (such as the $(k-1)$-sphere, which is usually embedded in $\mathbb{R}^k$), we may also refer to this measure as the *area* of the object. The meaning of these terms should be clear from context, and the Vol() and Area() predicates may be used.

The volume of $B_d(r)$ will be written as $V_d(r)$. It is a standard result [33] that

$$V_d(r) = \frac{\pi^{\frac{d}{2}} r^d}{\Gamma\left(\frac{d}{2} + 1\right)},$$

where $\Gamma(\cdot)$ denotes the (complete) gamma function.

For fixed r this quantity diminishes to zero as d goes to infinity, and $V_d(1)$ is bounded by $8\pi^2/15$ for all d. Also,

$$\mathrm{Area}(S_{d-1}(r)) \;=\; \frac{d\ V_d(r)}{r}.$$

Throughout the paper, the uppercase letter K, with or without a subscript or superscript, will always denote some constant value.

The model. Let the robot have a d-dimensional configuration space $\mathcal{C}$, defined by a polytope of unit diameter in $\mathbb{R}^d$. (The restriction on the diameter will be removed later, when we present the main theorem.) $\mathcal{C}$ is the domain from which the milestones are sampled by the PRM algorithm. Let D_{in} be the diameter of the largest ball contained completely within $\mathcal{C}$. The dimension d of the space and the domain parameter D_{in} will be considered constants in our treatment. Let $\mathcal{S}$ be a set of n $(d-1)$-simplices in $\mathcal{C}$. These are the configuration space obstacles in our model. Thus $\mathcal{C}_{\text{free}} = \mathcal{C} \setminus \bigcup_{s \in \mathcal{S}} s$.

A probability distribution $\mathcal{D}$ on $\mathbb{R}^d$ with density function $\mu(\cdot)$ is said to be γ-smooth, for some $\gamma \in \mathbb{R}$, if

1. $\mu(x) \le \gamma$ for all $x \in \mathbb{R}^d$, and
2. given any hyperplane H in $\mathbb{R}^d$, a point distributed under $\mathcal{D}$ is on H with probability 0.

A symmetric d-variate Gaussian distribution with variance σ^2 is $\Theta\left(\frac{1}{\sigma^d}\right)$-smooth. We assume that each vertex of each simplex in $\mathcal{S}$ is independently perturbed according to a γ-smooth distribution within the domain.

The locally orthogonal decomposition. The locally orthogonal decomposition $\Xi(\mathcal{S})$ of $\mathcal{S}$ is the arrangement of the following two collections of hyperplanes:

- The affine hull $\mathrm{Aff}(s)$ of s, for each $s \in \mathcal{S}$.
- The hyperplane orthogonal to s and containing f, for each $s \in \mathcal{S}$ and each facet f of s.

Hyperplanes of the second type are called *walls*. A facet of $\Xi(\mathcal{S})$ is *bound* if it is contained in some $s \in \mathcal{S}$, otherwise it is *free*. In the following, the decomposition is assumed to be restricted to $\mathcal{C}$. The second property of γ-smooth distributions ensures that under our perturbation model, $\Xi(\mathcal{S})$ is in general position with probability 1.

Lemma 1. *Let c_1 and c_2 be two cells of $\Xi(\mathcal{S})$ that are incident at a free facet. Then for any $p_1 \in c_1$ and $p_2 \in c_2$, the line segment between p_1 and p_2 is disjoint from $\mathcal{S}$.*

Proof. Let H be the hyperplane containing the facet that separates c_1 and c_2. H is part of $\Xi(\{s\})$ for some $s \in \mathcal{S}$. Let $\Xi(\{s\}) - H$ refer to the subdivision induced by the simplex s and all the hyperplanes of $\Xi(\{s\})$ other than H. $\Xi(\{s\}) - H$ is a convex subdivision: if H is the affine hull of s we have a prism split in half by s, otherwise we have a subdivision induced by a set of hyperplanes. Thus the overlay $\mathcal{O}$ of $\Xi(\mathcal{S} - \{s\})$ with $\Xi(\{s\}) - H$ is also a convex subdivision. The cells c_1 and c_2 lie in the same cell of $\mathcal{O}$. This implies the lemma. □

Corollary 2. *If a milestone is placed in each cell of $\Xi(\mathcal{S})$ then any two points that can be connected by a path in $\mathcal{C}_{free}$ can also be connected by a piecewise linear path whose only internal vertices are milestones.*

Proof. Let p and q be points in $\mathcal{C}_{\mathrm{free}}$ that can be connected by a feasible path Π. Let $\{c_1, c_2, \ldots, c_k\}$ be the sequence of cells of $\Xi(\mathcal{S})$ traversed by Π, and let m_i be a milestone in c_i. By Lemma 1, the piecewise linear path with vertices $\{p, m_1, m_2, \ldots, m_k, q\}$ is feasible. Figure 1 illustrates this. □

Volume bound. Corollary 2 implies that it suffices to place a milestone in every cell of $\Xi(\mathcal{S})$. To show that this can be accomplished with a polynomial number of samples we prove a high-probability lower bound on the volume of each cell of $\Xi(\mathcal{S})$. This is achieved with the help of the following simple lemma, which is easily proved by induction.

Lemma 3. *Let $\mathcal{A}(\mathcal{H})$ be the arrangement of a set of hyperplanes $\mathcal{H}$. If every vertex v of $\mathcal{A}(\mathcal{H})$ is ε-distant from every hyperplane $H \in \mathcal{H}$ for which $v \notin H$, then the volume (k-dimensional measure) of any k-face of the arrangement is at least $\varepsilon^k/k!$, for $1 \leq k \leq d$.*

This lemma implies that volume bounds can be proved through vertex-hyperplane separation bounds. Accordingly, Section 3 is devoted to proving the following theorem:

Theorem 4. *Consider a vertex v and a hyperplane H of $\Xi(\mathcal{S})$ such that $v \notin H$, and let $\Delta := \min\{1, D_{\mathrm{in}}\}$. Given $\varepsilon \in [0, \Delta)$, v is ε-close to H with probability at most*

$$O\left(\varepsilon^{1-\alpha} \max\{\gamma, \gamma^{d^2}\}\right)$$

for any $\alpha > 0$.

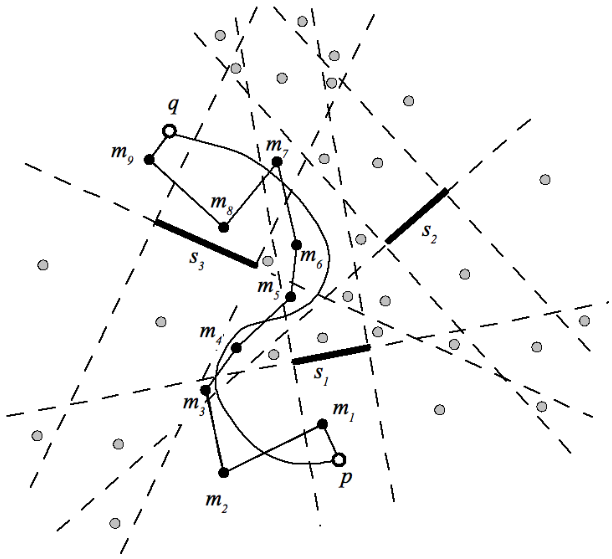

Figure 1: If two points p and q can be connected by a path in $\mathcal{C}_{\text{free}}$ they can also be connected by a linear interpolation of milestones $\{m_i\}$, as long as one is placed in each cell of the locally orthogonal decomposition.

The number of hyperplanes in $\Xi(\mathcal{S})$ is $O(n)$ and the number of vertices is $O(n^d)$. A union bound and an application of Lemma 3 thus yield the following corollary to Theorem 4.

Corollary 5. *Given* $\varepsilon \in \left[0, \frac{\Delta^d}{d!}\right)$, *the probability that some cell of* $\Xi(\mathcal{S})$ *has volume less than or equal to* ε *is*

$$O\left(n^{d+1}\varepsilon^{\frac{1-\alpha}{d}} \max\{\gamma, \gamma^{d^2}\}\right)$$

for any $\alpha > 0$. *Hence each cell has volume at least* ε *with probability at least* $1-\omega$ *if*

$$\varepsilon \;\le\; \min\left\{K\,\omega^{\frac{d}{1-\alpha}}\,n^{-\frac{d(d+1)}{1-\alpha}}\left(\max\{\gamma,\gamma^{d^2}\}\right)^{-\frac{d}{1-\alpha}}, \frac{\Delta^d}{d!}\right\}$$

for an appropriate constant K.

If each cell of $\Xi(\mathcal{S})$ has volume at least ε, standard probability theory implies that the expected number of samples sufficient for placing a milestone in every cell is $O\left(\frac{1}{\varepsilon}\log\frac{1}{\varepsilon}\right)$ [31]. Applying Corollary 5, we conclude that with high probability, a set of $\mathrm{Poly}(n,\gamma)$ samples from $\mathcal{C}$ is expected to place a milestone in every cell of $\Xi(\mathcal{S})$. This yields our main theorem, which we state for *arbitrarily large* domains in the special case of Gaussian perturbations.

Theorem 6. *For constant* d, *let a free configuration space be defined by* n $(d-1)$-*simplices in* $\mathbb{R}^d$ *within a polyhedral domain of diameter* D. *If independent Gaussian perturbations of variance* $(\sigma D)^2$ *are applied to the simplex vertices then the expected* *number of uniformly chosen random samples required to construct an accurate probabilistic roadmap is polynomial in n and $\frac{1}{\sigma}$.*

In the statement of this theorem, we have removed the restriction that the domain have unit diameter. For a domain of diameter $D \neq 1$, smoothed analysis requires us to apply a perturbation of variance $(\sigma D)^2$ (recall that the perturbation is proportional to some measure of the numerical magnitude of the input: the diameter is a good fit). When we scale the domain to unit diameter, the variance becomes σ^2 to maintain scale invariance of the problem. This is precisely the situation we studied for unit diameter, and the same bounds apply.

3. Distance Bounds

This section is devoted to proving Theorem 4, which upper-bounds the probability that a vertex v and a hyperplane H of $\Xi(\mathcal{S})$ are ε-close. The one-dimensional case admits a simple proof which does not require the decomposition machinery, so we assume $d \geq 2$ in the balance of this paper. H can fall into three categories:

1. The affine span of $s \in \mathcal{S}$.
2. A wall containing a facet of $s \in \mathcal{S}$.
3. A hyperplane defining the boundary of $\mathcal{C}$.

We analyze these cases separately, devoting a subsection to each.

3.1. Affine Spans of Simplices

Theorem 7. *Consider a fixed point p in $\mathbb{R}^d$. Given $0 \leq k < d$, let the points $U = \{u_1, u_2, \ldots, u_{k+1}\}$ be distributed independently and γ-smoothly in $\mathcal{C}$. The probability that the affine span of U is ε-close to p is at most*

$$K\varepsilon^{d-k}\gamma^{k+1}$$

for $\varepsilon \geq 0$ and a constant K that depends only on k and d.

Proof. For $k = 0$ the result is trivial. Assume $1 \leq k \leq \frac{d}{2}$. We will integrate over all k-flats formed by the affine span of $(k+1)$-tuples of points. Let F denote the affine span of U. For a given u_1, the k-subspace $F - u_1$ of $\mathbb{R}^d$ can be represented as the span of k orthonormal vectors $v_1, v_2, \ldots, v_k$. The constraint that the vectors must be orthonormal makes direct parametrization and integration difficult. Instead, we will show that there is a mapping from *arbitrary* k-tuples of unit vectors in $(d-k+1)$-space to orthonormal bases for k-subspaces of d-space that satisfies certain necessary properties. With this mapping in hand, we can integrate in a straightforward fashion over the former space instead.

Let N be an arbitrary k-tuple of unit vectors in $\mathbb{R}^{d-k+1}$, i.e. each vector is drawn from the $(d-k)$-dimensional unit sphere S_{d-k}. Assume for the moment that we have a mapping ϕ that maps N to an orthonormal basis for a k-subspace of $\mathbb{R}^d$ and satisfies the following properties:

- It is surjective, or onto, in the sense that every k-subspace of $\mathbb{R}^d$ has an orthonormal basis W such that there is some k-tuple N drawn from S_{d-k} with $\phi(N) = W$. Note that we do *not* require that every set of k orthonormal d-vectors have a pre-image under ϕ.

- It is continuous under a particular metric ρ_k. For any two k-tuples of vectors $W = (v_1, v_2, \ldots, v_k)$ and $W' = (v'_1, v'_2, \ldots, v'_k)$ in some space, define $\rho_k(W, W') = \sup_{i=1}^k \|v_i - v'_i\|$. We require that if N and N' are two k-tuples of unit vectors in $\mathbb{R}^{d-k+1}$ such that $\rho_k(N, N') \leq \delta$, then $\rho_k(\phi(N), \phi(N')) \leq K^*\delta$ for some constant K^* that depends only on k.

Now we divide S_{d-k} into differential elements $A_1, A_2, \ldots, A_m$. Assume that the subdivision scheme has the following properties. ($\mathrm{Diam}(A_i)$ and $\mathrm{Area}(A_i)$ denote the diameter (measured in $\mathbb{R}^{d-k+1}$) and area of A_i, respectively.)

Property 1: $\inf_i \mathrm{Area}(A_i) \;\geq\; C \sup_i \mathrm{Diam}(A_i)^{d-k}$ for all $i := 1 \ldots m$ and a positive constant C independent of m. That is, the differential elements are "round".

Property 2: $\sup_i \mathrm{Area}(A_i) \to 0$ as m increases.

It is not hard to prove that such a scheme exists: consider, for example, drawing an uniform grid on the surface of the cube inscribed in the sphere and radially projecting the grid onto the sphere. We omit the formal proof in our presentation.

Let $\delta = \sup_i \mathrm{Diam}(A_i)$. We now choose a representative point $\hat{\mathrm{n}}_i^0$ in each A_i. Given an index k-tuple $I := (i_1, \ldots, i_k)$, let N_I^0 denote the k-tuple $(\hat{\mathrm{n}}_{i_1}^0, \hat{\mathrm{n}}_{i_2}^0, \ldots, \hat{\mathrm{n}}_{i_k}^0)$ and write $\phi(N_I^0) := (v_1^0, \ldots, v_k^0)$. Let N_I denote a k-tuple of the form $(\hat{\mathrm{n}}_{i_1}, \ldots, \hat{\mathrm{n}}_{i_k})$, where each $\hat{\mathrm{n}}_{i_j} \in A_{i_j}$. Write $\phi(N_I) := (v_1, \ldots, v_k)$. By our continuity criterion above, $\|v_i^0 - v_i\| \leq K^*\delta$ for $1 \leq i \leq k$.

Let $q := \alpha_1 v_1 + \alpha_2 v_2 + \cdots + \alpha_k v_k$ be any point on $\mathrm{Span}(\phi(N_I))$ within unit distance from the origin: this implies that each coefficient α_i has absolute value at most 1. This is a crucial observation that relies on the orthonormality of the set $\phi(N_I)$. The "neighbour" of q on $\mathrm{Span}(\phi(N_I^0))$ is the point $q^0 := \alpha_1 v_1^0 + \alpha_2 v_2^0 + \cdots + \alpha_k v_k^0$. Now

$$\begin{aligned} \|q - q^0\| &= \|\alpha_1(v_1 - v_1^0) + \cdots + \alpha_k(v_k - v_k^0)\| \\ &\leq \|\alpha_1(v_1 - v_1^0)\| + \cdots + \|\alpha_k(v_k - v_k^0)\| \\ &\leq kK^*\delta. \end{aligned}$$

So every point on $\mathrm{Span}(\phi(N_I))$ within $\mathcal{C} - u_1$ is $O(\delta)$-close to $\mathrm{Span}(\phi(N_I^0))$. Now we can write

$$\Pr[F \text{ is } \varepsilon\text{-close to } p] \;\leq \sum_{I \in \{1,\ldots,m\}^k} \Pr[A \text{ and } B]$$

where $A :=$ "F is ε-close to p", and $B :=$ "$F - u_1 = \mathrm{Span}(\phi(N_I))$ for some N_I". The inequality results from the observation that while the mapping ϕ is onto, it is not one-to-one: a basis is unchanged if we permute its members, so many k-tuples drawn from S_{d-k} (which are ordered) map to the same (unordered) basis.

Write $F^0 := \mathrm{Span}(\phi(N_I^0))$. If B is satisfied then, within $\mathcal{C}$, F must be contained in the set of points $(kK^*\delta)$-close to $u_1 + F^0$. Let G be the ball of radius $kK^*\delta$ in the linear space $F^\perp$ orthogonal to F^0. Then the required region is $G \oplus (u_1 + F^0)$, which has volume (within $\mathcal{C}$) at most $V_{d-k}(kK^*\delta)V_k(1) = K'\delta^{d-k}$. Each of $u_2, \ldots, u_{k+1}$ must lie within this region, so $\Pr[B] \le (\gamma K'\delta^{d-k})^k$.

Assume $\delta \ll \varepsilon$. Let G_ε be the ball of radius $\varepsilon + kK^*\delta \approx \varepsilon$ in $F^\perp$. $\Pr[A \mid B]$ is 1 only if u_1 is in $G_\varepsilon \oplus (p + F^0)$, which has volume $K''\varepsilon^{d-k}$ within $\mathcal{C}$, and is 0 otherwise. So integrating the indicator function over all possible locations of u_1, we get $\Pr[A \mid B] \le \gamma K''\varepsilon^{d-k}$. Multiplying and applying Property 1:

$$\begin{aligned}\Pr[A \text{ and } B] &= \Pr[A \mid B]\Pr[B] \\ &\le K'''\varepsilon^{d-k}\gamma^{k+1}\delta^{k(d-k)} \\ &\le \frac{K'''}{C^k}\varepsilon^{d-k}\gamma^{k+1}\prod_{j=1}^{k}\text{Area}(A_{i_j})\end{aligned}$$

Summing over all possible k-tuples of indices, we have

$$\begin{aligned}\Pr\left[F \text{ is } \varepsilon\text{-close to } p\right] &\le \frac{K'''}{C^k}\gamma^{k+1}\varepsilon^{d-k}\text{Area}(S_{d-k})^k \\ &= K\varepsilon^{d-k}\gamma^{k+1}\end{aligned}$$

for a constant K that depends only on d and k.

To wrap things up, we must handle the case $k > \frac{d}{2}$. Observe that if F is a k-flat for such k, then the orthogonal complement of $F - u_1$ is a $(d-k)$-subspace which can be studied as above. Further, if two $(d-k)$-subspaces are defined by orthonormal bases that are pairwise $O(\delta)$-close, then their orthogonal complements must be $O(\delta)$-close within $\mathcal{C} - u_1$ (i.e. every point on one is $O(\delta)$-close to the other). Running through the above argument in this scenario yields an identical result.

All we need to do now is to construct an appropriate mapping ϕ. Let $N = (\hat{n}_1, \hat{n}_2, \ldots, \hat{n}_k)$ be a k-tuple of points drawn from S_{d-k}. We will start by rigidly embedding S_{d-k} in a canonical $(d-k+1)$-subspace of $\mathbb{R}^d$ with center at the origin. Let T_1 be the rigid transformation that achieves this. Now consider another $(d-k+1)$-subspace orthogonal to $T_1(\hat{n}_1)$ and similarly embed S_{d-k} in it with a rigid transformation $T_2^{\hat{n}_1}$. Now take a third subspace orthogonal to both $T_1(\hat{n}_1)$ and $T_2^{\hat{n}_1}(\hat{n}_2)$, embed S_{d-k} in it, and recurse in this way until we have considered k subspaces. We then have the mapping $\phi(N) = (T_1(\hat{n}_1), T_2^{\hat{n}_1}(\hat{n}_2), \ldots, T_k^{\hat{n}_1,\ldots,\hat{n}_{k-1}}(\hat{n}_k))$.

Consider an arbitrary k-subspace F_k of $\mathbb{R}^d$. It must intersect the first embedded sphere at at least one point v_1. Since $T_1(\hat{n}_1)$ covers every point on this sphere as $\hat{n}_1$ varies over S_{d-k}, $v_1 = T_1(\hat{n}_1)$ for some choice of $\hat{n}_1$. F_k also intersects the second sphere, embedded in a subspace orthogonal to v_1. Assume v_2 is a point of intersection. By similar reasoning, $v_2 = T_2^{\hat{n}_1}(\hat{n}_2)$ for some $\hat{n}_2$. Continue in this way until we have the tuple $(v_1, v_2, \ldots, v_k) := \phi((\hat{n}_1, \hat{n}_2, \ldots, \hat{n}_k))$: it is easily verified that this is an orthonormal basis for F_k. Hence ϕ is surjective.

It is more difficult to choose the transformations so that ϕ is continuous. We will make do with an intermediate result. Informally, we will take a known mapping for a single k-tuple and "fudge" it to obtain the mapping for other k-tuples "near" the first. By doing this for a large number of "known" k-tuples, we cover the domain of ϕ. The resulting mapping is not guaranteed to be continuous as we defined the term, but it turns out to be sufficient for our needs.

Assume some arbitrarily chosen valid sequence of transformations $T_1, T_2^{\hat{n}_1}, \ldots, T_k^{\hat{n}_1,\ldots,\hat{n}_{k-1}}$ for the k-tuple $N := (\hat{n}_1, \hat{n}_2, \ldots, \hat{n}_k)$. Consider another k-tuple $N' := (\hat{n}'_1, \hat{n}'_2, \ldots, \hat{n}'_k)$ such that $\rho_k(N, N') \leq \delta$. Associate with this k-tuple the following sequence of transformations:

$$\begin{aligned}
&T_1 \\
T_2^{\hat{n}'_1} &:= R_2 \circ T_2^{\hat{n}_1} \\
T_3^{\hat{n}'_1,\hat{n}'_2} &:= R_3 \circ T_3^{\hat{n}_1,\hat{n}_2} \\
&\vdots \\
T_k^{\hat{n}'_1,\ldots,\hat{n}'_{k-1}} &:= R_k \circ T_k^{\hat{n}_1,\ldots,\hat{n}_{k-1}}
\end{aligned}$$

Here, the R_i's are a set of rotations. R_2 maps $T_1(\hat{n}_1)$ to $T_1(\hat{n}'_1)$, R_3 maps $T_1(\hat{n}_1)$ to $T_1(\hat{n}'_1)$ and $T_2^{\hat{n}_1}(\hat{n}_2)$ to $T_2^{\hat{n}'_1}(\hat{n}'_2)$, and so on for longer and longer prefixes of the bases. Note that for any $\hat{n}$,

$$\begin{aligned}
\langle T_1(\hat{n}'_1), T_2^{\hat{n}'_1}(\hat{n})\rangle &= \langle (R_2 \circ T_1)(\hat{n}_1), (R_2 \circ T_2^{\hat{n}_1})(\hat{n})\rangle \\
&= \langle T_1(\hat{n}_1), T_2^{\hat{n}_1}(\hat{n})\rangle \quad \text{(rotation is orthogonal and preserves inner products by definition)} \\
&= 0 \quad \text{(by the definition of } T_2^{\hat{n}_1}\text{)}
\end{aligned}$$

Hence, our construction of $T_2^{\hat{n}'_1}$ is valid in that it maps S_{d-k} to a subspace orthogonal to $T_1(\hat{n}'_1)$. Similar arguments establish the validity of the other transformations.

We require each rotation to displace unit vectors in $\mathbb{R}^d$ by at most $K^{**}\delta$, where K^{**} depends only on k. To prove that such a sequence of rotations can be constructed, we state the following lemma.

Lemma 8. *Given $\delta \geq 0$ and $t \leq d$, let $(v_1, v_2, \ldots, v_t)$ and $(v'_1, v'_2, \ldots, v'_t)$ be two t-tuples of orthonormal vectors in $\mathbb{R}^d$, such that $\|v_i - v'_i\| \leq \delta$ for $1 \leq i \leq t$. Then there exists a rotation R of $\mathbb{R}^d$ about the origin that maps v_i to v'_i for $1 \leq i \leq t$, and maps each unit vector u in $\mathbb{R}^d$ to another unit vector u' such that $\|u - u'\| \leq K\delta$, where K depends only on t. Such a transformation R is called a $(K\delta)$-*rotation*.*

Proof. We will prove the result by induction on t. In the base case $t = 1$, if $v_1 = v'_1$ (which, incidentally, we have by default when $t = d = 1$) then R can be taken to be the identity transformation and we are done. Else, consider the 2-space H spanned by v_1 and v'_1. Without loss of generality, assume this plane is spanned by the first two basis vectors $\hat{x}_1$ and $\hat{x}_2$ in some canonical orthonormal basis B of $\mathbb{R}^d$. The two-dimensional rotation that maps v_1 to v'_1 in H, expressed in terms of the x_1 and x_2 components of a vector, is a standard result. Extend this rotation to d dimensions by stipulating that the final transformation R does not change any of the other components, i.e. the d-dimensional matrix for the transformation becomes identity when the first two rows and columns (corresponding to the x_1 and x_2 components) are deleted. Since this rotation R changes only the x_1 and x_2 components of a vector u when mapping it to u', the distance $\|u - u'\|$ is precisely the distance between the orthogonal projections of u and u' on H.

Call these projections u_H and u'_H respectively. It is straightforward to see that u'_H must be the image of u_H under the rotation R (in the basis B, they have the same first two components as u' and u respectively, and their other components are identical, being zero). Since R behaves as a two-dimensional rotation in H that maps v_1 to v'_1, with $\|v_1 - v'_1\| \leq \delta$, it must be that the distance between *any* vector of length at most one in H and its image under rotation is at most δ. In particular, this holds for u_H and its image u'_H, both of which have length at most $\|u\| = \|u'\| = 1$. By the reasoning above, we have $\|u - u'\| = \|u_H - u'_H\| \leq \delta$.

Now assume $t > 1$, and that the result holds for all $t', 1 \leq t' < t$. By the same reasoning as for the base case, there is a rotation R_t that maps v_1 to v'_1 and displaces any unit vector in $\mathbb{R}^d$ by a distance of at most δ. In particular, the vectors $v_2, v_3, \ldots, v_t$ move to new positions $v''_2, v''_3, \ldots, v''_t$ respectively such that each of $\|v_2 - v''_2\|, \|v_3 - v''_3\|, \ldots, \|v_t - v''_t\|$ is at most δ. By the triangle inequality, $\|v''_i - v'_i\| \leq \|v''_i - v_i\| + \|v_i - v'_i\| \leq 2\delta$, for $2 \leq i \leq t$. Note also that since R_t is a rotation, $(v'_1, v''_2, v''_3, \ldots, v''_t)$ is an orthonormal set. Apply the induction hypothesis in the $(d-1)$-dimensional orthogonal complement of v'_1, with the two $(t-1)$-tuples $(v''_2, v''_3, \ldots, v''_t)$ and $(v'_2, v'_3, \ldots, v'_t)$ which have pairwise distances bounded by 2δ: there is a rotation R^{d-1}_{t-1} in this subspace that maps the first tuple to the second and displaces any unit vector by a distance of at most $2K_{t-1}\delta$ for some K_{t-1} that depends only on t. Extend R^{d-1}_{t-1} to a d-dimensional rotation R_{t-1} by stipulating that the component of a d-vector along v'_1 remains unchanged: note in particular that R_{t-1} does not displace v'_1 itself. As in the base case, this extension ensures that R_{t-1} displaces any unit vector in $\mathbb{R}^d$ by at most $2K_{t-1}\delta$. Let $R = R_{t-1} \circ R_t$: this rotation maps $(v_1, v_2, \ldots, v_t)$ to $(v'_1, v'_2, \ldots, v'_t)$. By the triangle inequality again, R moves any unit vector in $\mathbb{R}^d$ by at most $(1 + 2K_{t-1})\delta = K\delta$ for some K that depends only on t, and is the required rotation. The result is thus proved by induction. □

Now let us return to the original discussion. Since T_1 is a rigid transformation, $\|T_1(\hat{n}_1) - T_1(\hat{n}'_1)\| = \|\hat{n}_1 - \hat{n}'_1\| \leq \delta$. Let R_2 be a $(K_2\delta)$-rotation, with K_2 depending only on k: its existence is guaranteed by Lemma 8. We have

$$\begin{aligned}
\|T_2^{\hat{n}'_1}(\hat{n}'_2) - T_2^{\hat{n}_1}(\hat{n}_2)\| &= \|(R_2 \circ T_2^{\hat{n}_1})(\hat{n}'_2) - T_2^{\hat{n}_1}(\hat{n}_2)\| \\
&\leq \|(R_2 \circ T_2^{\hat{n}_1})(\hat{n}'_2) - T_2^{\hat{n}_1}(\hat{n}'_2)\| + \|T_2^{\hat{n}_1}(\hat{n}'_2) - T_2^{\hat{n}_1}(\hat{n}_2)\| \\
&\leq K_2\delta + \delta \\
&= (K_2 + 1)\delta
\end{aligned}$$

Now we can recurse. Invoking Lemma 8 with $(K_2 + 1)\delta$ instead of δ and by using identical arguments, we establish that R_3 can be a $(K_3\delta)$-rotation, and similarly for $R_4, \ldots, R_k$. K^{**} is then the maximum of $K_2, K_3, \ldots, K_k$. Denote by ψ_N this restriction of ϕ to a domain of k-tuples δ-close to N. Observe that $\rho_k(\psi_N(N), \psi_N(N')) \leq K^*\delta$, where $K^* = K^{**} + 1$.

Recall that in our integration, we divided the sphere S_{d-k} into differential elements $A_1, A_2, \ldots, A_m$ of diameter δ and picked a representative point $\hat{n}_i^0$ from each A_i. We will define ϕ differently for each m: properly, we should replace it with a symbol such as ϕ_m, but for clarity of exposition we will be a bit loose with our notation. As before, given an index k-tuple $I := (i_1, \ldots, i_k)$, let N_I^0 denote the k-tuple $(\hat{n}^0_{i_1}, \hat{n}^0_{i_2}, \ldots, \hat{n}^0_{i_k})$, and let N_I denote a k-tuple of the form $(\hat{n}_{i_1}, \ldots, \hat{n}_{i_k})$, where each $\hat{n}_{i_j} \in A_{i_j}$. Define

$\phi(N_I) := \psi_{N_I^0}(N_I)$: thus we are guaranteed that $\rho_k(\phi(N_I), \phi(N_I^0)) \leq K^*\delta$. In other words, the mapping ϕ is continuous within each unique sequence of k differential elements. Since the main proof only compares the image of a k-tuple of points from a sequence of differential elements to that of its representative k-tuple from the same sequence, this "local" version of continuity is enough for the proof to go through. □

We derive the following corollary of Theorem 7.

Corollary 9. *For nonnegative integers k, k' that satisfy $k+k' < d$, consider an arbitrarily distributed k'-flat F in $\mathbb{R}^d$, as well as a set $U = \{u_1, \ldots, u_{k+1}\}$ of γ-smoothly distributed points in $\mathcal{C}$, independent of F and of each other. The shortest distance between F and Aff(U) is at most ε with probability at most*

$$K\varepsilon^{d-k-k'}\gamma^{k+1}$$

for $\varepsilon \geq 0$ and a constant K that depends only on k and d.

Proof. For $k' = 0$, Theorem 7 immediately yields the result: since the point F is distributed independently of U, we can hold it fixed, apply the theorem and then integrate the result over the range of F — it is trivially verified that this last step does not change the probability bound from the previous step. For $k' > 0$, fix F: by independence, the points in U retain their original distributions under this restriction. Let F_0 be the subspace of $\mathbb{R}^d$ identical to F except for translation, and let $F^\perp$ be the orthogonal complement of F_0. Evidently, the shortest distance of a point to F is preserved under orthogonal projection to $F^\perp$. F itself maps to a single point p of $F^\perp$. Further, the points in $\mathbb{R}^d$ mapping to a volume element $\mathrm{d}\sigma$ of $F^\perp$ are exactly those in $\mathrm{d}\sigma \oplus F_0$. The k'-area of any k'-flat, when restricted to $\mathcal{C}$, is at most $V_{k'}(1)$, so the volume of the Minkowski sum (in $\mathcal{C}$) is at most $V_{k'}(1)\mathrm{d}\sigma$. This is illustrated in Figure 2. Hence the projection $u_i^\perp$ of each u_i is $(\gamma V_{k'}(1))$-smoothly and independently distributed in $F^\perp$. Now we can apply Theorem 7 in the $(d-k')$-dimensional space $F^\perp$ to upper-bound the probability that p is ε-close to $\mathrm{Aff}(u_1^\perp, \ldots, u_{k+1}^\perp)$, and hence the probability that F is ε-close to Aff(U), by

$$K\varepsilon^{d-k-k'}\gamma^{k+1}$$

This has no dependence on F, so integrating over the distribution of F gives the same overall probability that F is ε-close to Aff(U). The formula simplifies to the required result. □

From Theorem 7 we see that a hyperplane-vertex pair of $\Xi(\mathcal{S})$, in which the hyperplane is the affine span of a simplex s, and the vertex v is defined entirely by hyperplanes not associated with s, is ε-close with probability at most polynomial in ε and γ. Specifically, the bound is $K\varepsilon\gamma^d$ for a constant K that depends only on d.

Corollary 9 applies to the case when the vertex is formed by the intersection of one or more walls supporting s with hyperplanes not associated with s. We extend the use of the term "wall" as follows: The intersection of a number of walls of s is the wall W spanned by Aff(U) and the normal to s, for a subset U of the vertices of s. Since W is orthogonal to s and contains v, we have

$$\mathrm{dist}(\mathrm{Aff}(s), v) \;=\; \mathrm{dist}(\mathrm{Aff}(U), v) \;\geq\; \mathrm{dist}(\mathrm{Aff}(U), Z)$$

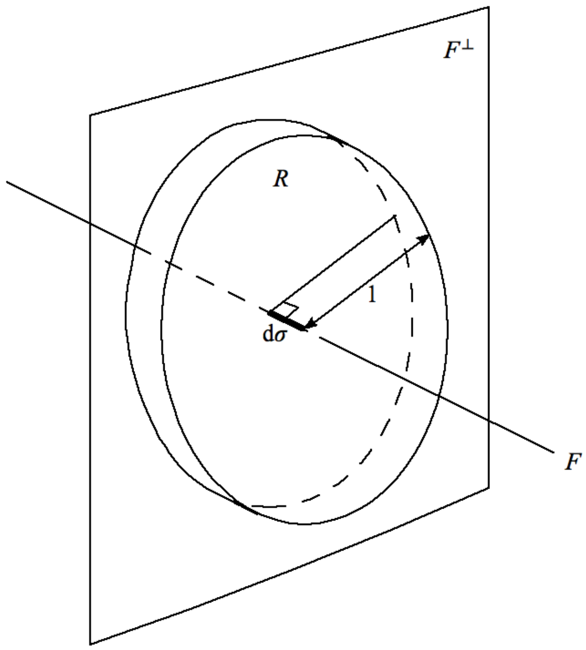

Figure 2: F is a (1D) subspace of $\mathbb{R}^3$, and $F^\perp$ is its (2D) orthogonal complement. dσ is a small volume (here, length) element of F and R is a ball of unit radius in $F^\perp$. d$\sigma \oplus R$ (here, a cylinder) contains all points within the domain that orthogonally project onto dσ.

where Z is the intersection of the hyperplanes unrelated to s. W and Z intersect at a point, so $\dim(W) + \dim(Z) = d$. Also, $\dim(\text{Aff}(U)) = \dim(W) - 1$, and the points in U are distributed γ-smoothly and independently (of each other and of Z) in $\mathcal{C}$. These are precisely the conditions required to apply Corollary 9, giving an upper bound on the probability that $\text{dist}(\text{Aff}(U), Z) \leq \varepsilon$, and hence on the probability that $\text{dist}(\text{Aff}(s), v) \leq \varepsilon$, that is again polynomial in ε and γ. Specifically, if k is the cardinality of U, then the bound is $K\varepsilon\gamma^k$ for a constant K that depends only on k and d.

3.2. Walls Supporting Simplices

When the hyperplane is a wall containing a simplex facet, the analysis is trickier. We will divide it into three cases based on the relationship between the wall and the vertex. These cases may be summarized as:

1. The wall and the vertex are independent.
2. The wall and the vertex depend on the same simplex but the vertex does not lie in the affine span of that simplex.
3. The wall and the vertex depend on the same simplex and the vertex lies in the affine span of that simplex.

Case 1. We will assume that the simplex associated with the wall is independent of the vertex and prove a rather general result.

Theorem 10. *Consider a simplex $s \in \mathcal{S}$. For nonnegative integers k, k' that satisfy $k + k' < d$, consider a subset $U = \{u_1, u_2, \ldots, u_k\}$ of the vertices of s and let W be the wall spanned by $Q := \mathit{Aff}(U)$ and the normal to s. Let F be a random k'-flat whose distribution is independent of s. The probability that W is ε-close to F is at most*

$$K\varepsilon^{d-k-k'}\gamma^{d-1}$$

for $\varepsilon \geq 0$ and a constant K that depends only on k and d.

Proof. Let H be the affine span of the simplex. Fix F, and let F^H be the orthogonal projection of F to H. By orthogonality, it is immediate that $\mathrm{dist}(W, F) = \mathrm{dist}(Q, F^H)$. We assume a tessellation scheme of S_{d-1} into area elements $A_1, A_2, \ldots, A_m$ as in the proof of Theorem 7, satisfying Properties 1 and 2. Write

$$\begin{aligned} \Pr[W \text{ is } \varepsilon\text{-close to } F] &= \sum_{i=1}^{m} \Pr[\hat{\mathrm{n}}(H) \in A_i \text{ and } W \text{ is } \varepsilon\text{-close to } F] \\ &= \sum_{i=1}^{m} \Pr\left[\hat{\mathrm{n}}(H) \in A_i \text{ and } Q \text{ is } \varepsilon\text{-close to } F^H\right] \end{aligned}$$

Now fix an arbitrary normal $\hat{\mathrm{n}}_i$ in each A_i and let H_0 be the plane $\langle x, \hat{\mathrm{n}}_i\rangle = 0$. Another normal $\hat{\mathrm{n}}$ also in A_i satisfies $\|\hat{\mathrm{n}} - \hat{\mathrm{n}}_i\| \leq \mathrm{Diam}(A_i)$. We will show that when $\hat{\mathrm{n}}(H)$, the unit normal to the hyperplane H, is in A_i, projection to H_0 instead of to H will almost surely not change the shortest distance by "much". For this the following simple lemma is required.

Lemma 11. *Let v_1 and v_2 be the orthogonal projections of vector v onto hyperplanes H_1 and H_2, respectively, and assume that the normals of H_1 and H_2 are δ-close, i.e., $\|\hat{\mathrm{n}}_2 - \hat{\mathrm{n}}_1\| \leq \delta$. Then $\|v_2\| - \|v_1\| \ \leq\ 2\delta\|v\|$.*

Proof. Write $\Delta n := \hat{\mathrm{n}}_2 - \hat{\mathrm{n}}_1$. We have $v_1 = v - \langle \hat{\mathrm{n}}_1, v\rangle\hat{\mathrm{n}}_1$ and $v_2 = v - \langle \hat{\mathrm{n}}_2, v\rangle\hat{\mathrm{n}}_2$. So

$$\begin{aligned} \|v_2\| - \|v_1\| &\leq \|v_2 - v_1\| \\ &\leq \|(v - \langle \hat{\mathrm{n}}_1, v\rangle\hat{\mathrm{n}}_1) - (v - \langle \hat{\mathrm{n}}_2, v\rangle\hat{\mathrm{n}}_2)\| \\ &= \|\langle \hat{\mathrm{n}}_2, v\rangle\hat{\mathrm{n}}_2 - \langle \hat{\mathrm{n}}_1, v\rangle\hat{\mathrm{n}}_1)\| \\ &= \|\langle \hat{\mathrm{n}}_1 + \Delta n, v\rangle(\hat{\mathrm{n}}_1 + \Delta n) - \langle \hat{\mathrm{n}}_1, v\rangle\hat{\mathrm{n}}_1)\| \\ &= \|\langle \hat{\mathrm{n}}_2, v\rangle\Delta n + \langle \Delta n, v\rangle\hat{\mathrm{n}}_1)\| \\ &\leq 2\|\Delta n\|\|v\| \ \leq\ 2\delta\|v\| \end{aligned}$$

□

Let (q, f^H) be a pair in $Q \times F^H$ such that $\mathrm{dist}(q, f^H) = \mathrm{dist}(Q, F^H)$ and let f be the pre-image of f^H under the projection—if there are multiple pre-images we choose the one closest to q. Let Q^{H_0}, F^{H_0}, q^{H_0} and f^{H_0} be the orthogonal projections of Q, F, q and f respectively to H_0. By the above lemma, $\mathrm{dist}(Q^{H_0}, F^{H_0}) \leq \mathrm{dist}(q^{H_0}, f^{H_0}) \leq$

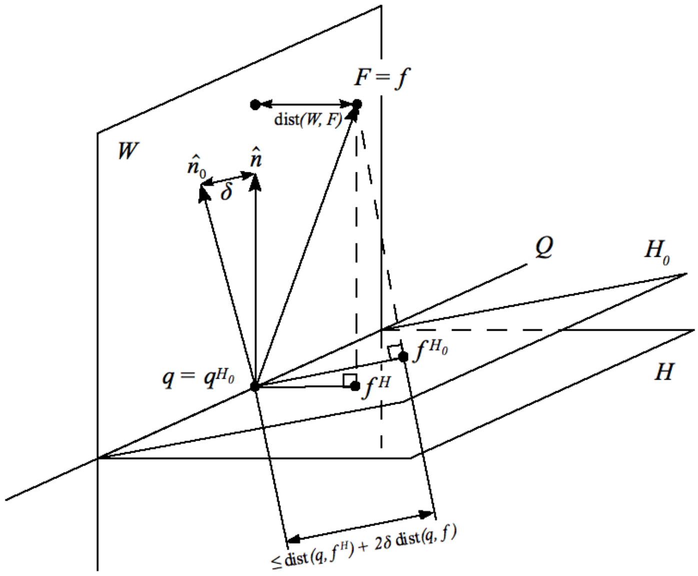

Figure 3: The projections of vector $\overrightarrow{qf}$ onto two hyperplanes H and H_0 differ by at most 2δ dist(q, f), where δ is the length of the difference of the normals of the hyperplanes. The notation is that of Theorem 10.

$\text{dist}(q, f^H) + 2\delta_i \text{dist}(q, f) = \text{dist}(Q, F^H) + 2\delta_i \text{dist}(q, f)$, where $\delta_i := \text{Diam}(A_i)$ (see Figure 3). For every possible configuration of s and F, $\text{dist}(q, f)$ is a finite positive quantity, so given $\omega \in (0, 1)$ we can always find a large enough constant M such that $\text{dist}(q, f) \leq M$ with probability at least ω. This implies

$$
\begin{aligned}
&\Pr\left[\hat{\text{n}}(H) \in A_i \text{ and } Q \text{ is } \varepsilon\text{-close to } F^H\right] \\
&\quad = \Pr[\hat{\text{n}}(H) \in A_i \text{ and } \text{dist}(q, f) \leq M \text{and } Q \text{ is } \varepsilon\text{-close to } F^H] \\
&\qquad\quad + \Pr[\hat{\text{n}}(H) \in A_i \text{ and } \text{dist}(q, f) > M \text{and } Q \text{ is } \varepsilon\text{-close to } F^H] \\
&\quad \leq \Pr[\hat{\text{n}}(H) \in A_i \text{ and } \text{dist}(q, f) \leq M \text{and } Q^{H_0} \text{ is } (\varepsilon + 2\delta_i M)\text{-close to } F^{H_0}] \\
&\qquad\quad + \Pr[\hat{\text{n}}(H) \in A_i \text{ and } \text{dist}(q, f) > M] \\
&\quad \leq \Pr[\hat{\text{n}}(H) \in A_i \text{and } Q^{H_0} \text{ is } (\varepsilon + 2\delta_i M)\text{-close to } F^{H_0}] \\
&\qquad\quad + \Pr[\hat{\text{n}}(H) \in A_i \text{ and } \text{dist}(q, f) > M]
\end{aligned}
$$

Let u_d be a vertex of s not in U. For the first term, observe that $\hat{\text{n}}(H)$ is in A_i only if every vertex of s is in the slab T between the parallel planes $\langle x - u_d, \hat{\text{n}}_i \rangle = \pm\delta_i$. The probability that a single vertex, considered independently from other factors, is in this slab is at most $2\gamma\delta_i V_{d-1}(1)$. Let $u_i^{H_0}$ be the orthogonal projection of each $u_i \in U$ on H_0 — under the above restriction, the $u_i^{H_0}$'s are $(2\gamma\delta_i)$-smoothly and independently distributed on H_0. Corollary 9 can now be directly applied in H_0 to obtain, for some constants K_1 and K_2,

$$
\begin{aligned}
&\Pr\left[\hat{\mathrm{n}}(H) \in A_i \text{ and } Q^{H_0} \text{ is } (\varepsilon + 2\delta_i M)\text{-close to } F^{H_0}\right] \\
&\quad \le \underbrace{(2\gamma\delta_i V_{d-1}(1))^{d-k-1}}_{\text{vertices not in } U \cup \{u_d\}} \times \underbrace{K_1(\varepsilon + 2\delta_i M)^{d-k-k'}(2\gamma\delta_i)^k}_{\text{vertices in } U} \\
&\quad \le K_2(\varepsilon + 2\delta_i M)^{d-k-k'}\gamma^{d-1}\mathrm{Area}(A_i).
\end{aligned}
$$

Summing over all i,

$$
\begin{aligned}
&\Pr[W \text{ is } \varepsilon\text{-close to } F] \\
&\quad \le \sum_i K_2(\varepsilon + 2\delta_i M)^{d-k-k'}\gamma^{d-1}\mathrm{Area}(A_i) \\
&\qquad\quad + \sum_i \Pr[\hat{\mathrm{n}}(H) \in A_i \text{ and } \mathrm{dist}(q, F) > M] \\
&\quad \le K_2(\varepsilon + 2\sup_i \delta_i M)^{d-k-k'}\gamma^{d-1}\mathrm{Area}(S_{d-1}) \\
&\qquad\quad + \Pr[\mathrm{dist}(q, F) > M] \\
&\quad \le \frac{8\pi^2 K_2}{15}(\varepsilon + 2\sup_i \delta_i M)^{d-k-k'}\gamma^{d-1} + (1-\omega)
\end{aligned}
$$

Make ω arbitrarily close to 1 and choose small enough area elements so that $\sup_i \delta_i M \ll \varepsilon$, thus obtaining

$$
\Pr[W \text{ is } \varepsilon\text{-close to } F] \;\le\; K\varepsilon^{d-k-k'}\gamma^{d-1}
$$

for a constant K. By independence, integrating over the range of F does not change the expression. □

Setting $k = d-1$ and $k' = 0$ yields the required vertex-wall separation result: The probability of ε-closeness is at most $K\varepsilon\gamma^{d-1}$.

Case 2. The next case to be treated is when the simplex associated with the wall is involved in the definition of the vertex but does not contain it in its affine span.

Theorem 12. *Consider a simplex $s \in \mathcal{S}$. Given a set $U = \{u_1, u_2, \ldots, u_k\}$ of k vertices of s, for $1 \le k < d$, define $Q := \mathit{Aff}(U)$. Let Z be the wall spanned by Q and the normal to s, let F be a random $(d-k)$-flat whose distribution is independent of s, and define $v := Z \cap F$. Let W be a wall of $\Xi(\{s\})$ that does not contain Z. Given $\varepsilon \in [0, 1)$, the probability that v is ε-close to W is at most*

$$
K\varepsilon^{1-\alpha}\gamma^{d-1}
$$

for any $\alpha > 0$ and a constant K that depends on α, k and d.

Proof. For $d \le 2$ the proof is straightforward. Assume $d > 2$. Let H be the affine span of the simplex. Assume, without loss of generality, that W contains $W_b := \mathrm{Aff}(u_2, u_3, \ldots, u_d)$. The intersection of W and Z is the wall Y spanned by $Y_b := \mathrm{Aff}(u_2, u_3, \ldots, u_k)$ and the normal to s (refer to Figure 4 for an illustration in three dimensions). Consider the

$(d-k+1)$-dimensional linear space $Y^{\perp}$ orthogonal to Y, in which Y itself orthogonally projects to a point y and Z projects to a line L. Note that $Y^{\perp}$ must contain $\hat{\mathrm{n}}(W)$. Let $\vec{a}$ be the vector $v^{\perp}-y$, where $v^{\perp}$ is the orthogonal projection of v to $Y^{\perp}$, lying on L. Since the projection to $Y^{\perp}$ is orthogonal to the perpendicular from v on W, it preserves the length of this vector, i.e. the perpendicular distance of v from W. Hence, we may measure this distance *after* projection as the dot product $|\langle\vec{a},\hat{\mathrm{n}}(W)\rangle|$. If v is ε-close to W, this implies that $\langle\vec{a},\hat{\mathrm{n}}(W)\rangle|\leq\varepsilon$. Figure 4 illustrates that $v^{\perp}$ must then be $(\varepsilon\sec\Theta)$-close to y, where Θ is the (acute) angle between L and $\hat{\mathrm{n}}(W)$. This means that $\mathrm{dist}(F,Y)\leq\mathrm{dist}(v,Y)=\mathrm{dist}(v^{\perp},y)\leq\varepsilon\sec\Theta$. Hence for infinitesimally small $\mathrm{d}\theta$,

$$\Pr\left[v \text{ is } \varepsilon\text{-close to } W \text{ and } \Theta\in[\theta,\theta+\mathrm{d}\theta)\right] \;\leq\; \Pr\left[F \text{ is } (\varepsilon\sec\theta)\text{-close to } Y \text{ and } \Theta\in[\theta,\theta+\mathrm{d}\theta)\right]$$

We now localize the normal of H and follow the proof of Theorem 10 with a few changes. Specifically, the probability

$$\Pr\left[\hat{\mathrm{n}}(H)\in A_i \text{ and } Q^{H_0} \text{ is } (\varepsilon+2\delta_i M)\text{-close to } F^{H_0}\right]$$

is replaced with the probability $P=\int_0^{\pi/2}P_\theta$, where

$$P_\theta \;:=\; \Pr[\hat{\mathrm{n}}(H)\in A_i \text{ and } \Theta\in[\theta,\theta+\mathrm{d}\theta)] \text{ and } Y_b^{H_0} \text{ is } ((\varepsilon+2\delta_i M)\sec\theta)\text{-close to } F^{H_0}]$$

(The differential $\mathrm{d}\theta$ will be shown to be present as a factor in P_θ.)

The fixed point is taken to be u_d as before. Note that if W_b is fixed then L, and hence the angle Θ, depends only on the position of u_1. In $Y^{\perp}$, let R be the region between the double cones with vertex y, axis $\hat{\mathrm{n}}(W)$ and half-angles θ and $\theta+\mathrm{d}\theta$. Evidently, Θ lies in the required range if and only if u_1 lies in the extruded region $R_Y:=R\oplus Y_b$. This is illustrated in Figure 4.

Even if W_b and hence Y_b are fixed, R depends on $Y^{\perp}$ and thus on $\hat{\mathrm{n}}(W)$ and on u_1. This yields a circularity. However, $\hat{\mathrm{n}}(W)$ can be approximated by a single normal $\hat{\mathrm{n}}_i$ in A_i — the approximation improves as A_i shrinks. This fixes Y independently of u_1 (denote this value by Y_0) and places R in $Y_0^{\perp}$ as the region R^0, which extrudes to R_Y^0 in H_0. The required probability can now be approximated as

$$\Pr\left[u_1 \text{ lies in } R_Y\right] \;\approx\; \Pr\left[u_1^{H_0} \text{ lies in } R_Y^0\right],$$

where $u_1^{H_0}$ is the orthogonal projection of u_1 to H_0. R^0, in domain H_0, has measure at most

$$\frac{2\mathrm{Area}(S_{d-k-1}(\sin\theta))\times 1^2\mathrm{d}\theta}{d-k+1} \;\leq\; \frac{2(d-k)V_{d-k}(1)\mathrm{d}\theta}{d-k+1}$$

as may be verified by picturing R^0 as a rotational sweep of a 2D double-cone. This implies that the volume of R_Y^0 within $\mathcal{C}$ is at most

$$\frac{1^{k-1}\;\times\;2(d-k)V_{d-k}(1)\;\mathrm{d}\theta}{d-k+1}.$$

We can now evaluate the required probability using

$$\begin{aligned}
P_\theta \;&\leq\; \Pr[\hat{\mathrm{n}}(H)\in A_i \text{ and } \Theta\in[\theta,\theta+\mathrm{d}\theta) \;\mid\; B_1]\\
&\quad\times\;\Pr[\hat{\mathrm{n}}(H)\in A_i \text{ and } Y_b^{H_0} \text{ is } ((\varepsilon+2\delta_i M)\sec\theta)\text{-close to } F^{H_0}]\\
&\leq\; \Pr[u_1\in T \text{ and } u_1^{H_0} \text{ is in } R_Y^0 \;\mid\; B_2]\\
&\quad\times\;\Pr[\{u_2,\ldots,u_{d-1}\}\subset T \text{and } Y_b^{H_0} \text{ is } ((\varepsilon+2\delta_i M)\sec\theta)\text{-close to } F^{H_0}]
\end{aligned}$$

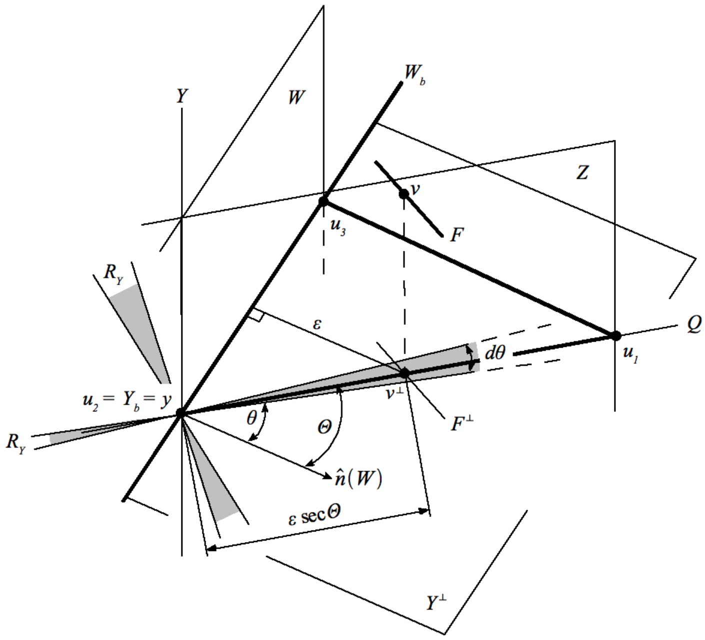

Figure 4: A vertex v is formed by the intersection of flat F and wall Z supporting the simplex $u_1u_2u_3$. Using the notation of Theorem 12, if L is at an angle $\Theta \in [\theta, \theta + \mathrm{d}\theta)$ to $\hat{\mathrm{n}}(W)$ in $Y^\perp$, and v is at distance ε from W, then the projection $v^\perp$ of v to $Y^\perp$ is at distance $\varepsilon \sec \Theta$ from y, the projection of Y to $Y^\perp$. L is at the required angle if and only if u_1 is in the shaded region R_Y.

where B_1 and B_2 are the conditions in the corresponding second factors in the two lines and T is the usual $2\delta_i$-thick slab for localizing the normal to the differential element. Note that the first factor in the last line depends only on u_1 and the second only on $u_2, \ldots, u_d$. The first factor is

$$\begin{aligned} &\Pr\left[u_1 \in T \text{ and } u_1^{H_0} \text{ is in } R_Y^0 \;\mid\; B_2\right] \\ &\quad\leq\; 2\gamma\delta_i \frac{2(d-k)\; V_{d-k}(1)\; \mathrm{d}\theta}{d-k+1} \\ &\quad=\; K_1\gamma\delta_i \mathrm{d}\theta \end{aligned}$$

for an appropriate constant K_1. The second factor is as in Theorem 10 (minus the vertex u_1 of U, and with an extra $\sec\theta$ factor), i.e. it is at most

$$K_2(\varepsilon + 2\delta_i M) \sec\theta \gamma^{d-2} \delta_i^{d-2},$$

for another constant K_2. Multiplying the bounds yields

$$\begin{aligned} P_\theta &\leq K_1 K_2 (\varepsilon + 2\delta_i M) \sec\theta \gamma^{d-1} \delta_i^{d-1} \mathrm{d}\theta \\ &\leq \frac{K_1 K_2}{C} (\varepsilon + 2\delta_i M) \gamma^{d-1} \mathrm{Area}(A_i) \sec\theta \mathrm{d}\theta \end{aligned}$$

The probability P is thus $\int_0^{\pi/2} P_\theta$, which is at most

$$\begin{aligned}
&\frac{K_1 K_2}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\text{Area}(A_i)\int_0^{\pi/2}\sec\theta \mathrm{d}\theta \\
&= \frac{K_1 K_2}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\text{Area}(A_i)[\log(\sec\theta + \tan\theta)]_0^{\pi/2}
\end{aligned}$$

Unfortunately this integral is unbounded. To circumvent this problem we write $P = P_a + P_b$, for

$$P_a = \int_{\pi/2-(\varepsilon+2\delta_i M)}^{\pi/2} P_\theta \quad \text{and} \quad P_b = \int_0^{\pi/2-(\varepsilon+2\delta_i M)} P_\theta$$

Then

$$\begin{aligned}
P_a &\le \int_{\pi/2-(\varepsilon+2\delta_i M)}^{\pi/2} \Pr\left[u_1 \in T \text{ and } u_1^{H_0} \text{ is in } R_Y^0 \mid B_2\right] \times \Pr\left[\{u_2, \ldots u_{d-1}\} \subset T\right] \\
&\le K_1\gamma\delta_i(\varepsilon + 2\delta_i M) \times K_3\gamma^{d-2}\delta_i^{d-2} \\
&\le \frac{K_1 K_3}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\text{Area}(A_i)
\end{aligned}$$

for some constant K_3, and

$$P_b \le \frac{K_1 K_2}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\text{Area}(A_i) \times [\log(\sec\theta + \tan\theta)]_0^{\pi/2-(\varepsilon+2\delta_i M)}$$

Now, for $0 < x < \pi/2$,

$$\begin{aligned}
\log(\sec\theta + \tan\theta)\mid_{\pi/2-x} &= \log(\operatorname{cosec} x + \cot x) \\
&= \log\frac{1+\cos x}{\sin x} \\
&\le \log\frac{2}{2x/\pi} \\
&= \log\frac{\pi}{x} \le K_\alpha\left(\frac{\pi}{x}\right)^\alpha
\end{aligned}$$

for any $\alpha > 0$ and a constant K_α that depends on α. Thus

$$\begin{aligned}
P_b &\le \frac{K_1 K_2}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\text{Area}(A_i)K_\alpha\left(\frac{\pi}{\varepsilon + 2\delta_i M}\right)^\alpha \\
&= \frac{K_1 K_2 K_\alpha \pi^\alpha}{C}(\varepsilon + 2\delta_i M)^{1-\alpha}\ \gamma^{d-1}\text{Area}(A_i)
\end{aligned}$$

The previous arguments imply that we can choose small enough area elements so that $\varepsilon + 2\sup_i \delta_i M \to \varepsilon < 1$. Therefore,

$$P = P_a + P_b \le K_4\varepsilon^{1-\alpha}\ \gamma^{d-1}\text{Area}(A_i)$$

for another constant K_4. Summing over all i,

$$\begin{aligned}
\Pr[v \text{ is } \varepsilon\text{-close to } W] &\le K_4\varepsilon^{1-\alpha}\ \gamma^{d-1}\ \text{Area}(S_{d-1}) \\
&\le \frac{8\pi^2 d K_4}{15}\varepsilon^{1-\alpha}\ \gamma^{d-1}
\end{aligned}$$

□

Case 3. The third case is when the affine span of the simplex associated with the wall is one of the hyperplanes defining the vertex.

Theorem 13. *Consider a simplex $s \in \mathcal{S}$. Given a set $U = \{u_1, u_2, \dots, u_k\}$ of k vertices of s, for $1 \leq k \leq d$, define $Q := \mathit{Aff}(U)$. Let F be a random $(d-k+1)$-flat whose distribution is independent of s, and define $v := Q \cap F$. Let W be a wall of $\Xi(\{s\})$ that does not contain Q. Given $\varepsilon \in [0,1)$, the probability that v is ε-close to W is at most*

$$K\varepsilon^{1-\alpha} \max\{\gamma, \gamma^k\}$$

for any $\alpha > 0$ and a constant K that depends on α, k and d.

Proof. The proof is similar to, but simpler than, that of Theorem 12. Because the intersection point lies in the affine span H of the simplex, the localization of the normal and subsequent projection onto this span is unnecessary. Assume, without loss of generality, that W stands on $W_b := \mathrm{Aff}(u_2, u_3, \dots, u_d)$. The intersection of W_b and Q is $Y := \mathrm{Aff}(u_2, u_3, \dots, u_k)$. Since the simplex and the wall are orthogonal, $\mathrm{dist}(W, v) = \mathrm{dist}(W_b, v)$ and we can restrict our attention to the hyperplane H: our next few comments will pertain strictly to this domain. Let $Y^\perp$ be the orthogonal complement of Y (w.r.t. H). The orthogonal projection of Y to $Y^\perp$ is the single point y, that of Q is the line L, and that of v is a point $v^\perp$ lying on L. Let the normal to W_b be $\hat{\mathrm{n}}(W_b)$. $\mathrm{dist}(W_b, v) = \mathrm{dist}(y, v^\perp)\cos\theta \geq \mathrm{dist}(Y, F)\cos\theta$, where θ is the measure of the angle Θ between $\hat{\mathrm{n}}(W_b)$ and $v^\perp - y$. Let R be the region between the double cones with vertex y, axis $\hat{\mathrm{n}}(W_b)$ and half-angles θ and $\theta + \mathrm{d}\theta$. Evidently, Θ lies in the required range iff u_1 lies in the extruded region $R_Y := R \oplus Y$, which has volume at most

$$\frac{2(d-k)\ V_{d-k}(1)\ \mathrm{d}\theta}{d-k+1}$$

Now we shift our attention back to the full-dimensional domain. Given $\{u_2, u_3, \dots, u_d\}$, Θ is in the required range iff u_1 lies in the region R' swept out by rotating R_Y around the axis W_b. By a rough estimate, the volume of this region is at most $2\pi \times 1 \times \mathrm{Vol}(R_Y)$, so

$$\Pr[\Theta \in [\theta, \theta + \mathrm{d}\theta)] \ = \ \gamma \mathrm{Vol}(R') \ = \ K_1 \gamma \mathrm{d}\theta$$

for a suitable constant K_1 that depends on k and d, and

$$\begin{aligned} P_\theta \ &:= \ \Pr[\Theta \in [\theta, \theta + \mathrm{d}\theta) \text{ and } Y \text{ is } (\varepsilon \sec\theta)\text{-close to } F] \\ &\leq \ K_1 \gamma \mathrm{d}\theta \times K_2 \varepsilon \sec\theta \gamma^{k-1} \end{aligned}$$

where K_2 is another constant that depends on k and d, from Corollary 9. Now as in the proof of Theorem 12, integrating this upper bound for P_θ over $[0, \pi/2]$ gives an unbounded result. We reuse our earlier hack to solve this problem. First,

$$\int_{\pi/2-\varepsilon}^{\pi/2} P_\theta \ \leq \ \int_{\pi/2-\varepsilon}^{\pi/2} \Pr[\Theta \in [\theta, \theta + \mathrm{d}\theta)] \ \leq \ K_1 \varepsilon \gamma$$

Next, for any $\alpha > 0$,

$$\int_0^{\pi/2-\varepsilon} P_\theta \ \leq \ K_1 K_2 K_\alpha \varepsilon^{1-\alpha} \gamma^k$$

for a constant K_α that depends on α. Putting everything together,

$$\int_0^{\pi/2} P_\theta \;\leq\; K\varepsilon^{1-\alpha}\max\{\gamma,\gamma^k\}$$

for a constant K that depends on α, k and d. □

3.3. The Boundary of the Domain

For this final case, we must bound the probability that a vertex v of $\Xi(\mathcal{S})$ not contained in a hyperplane H constituting the boundary of the domain $\mathcal{C}$ is ε-close to it. Since this hyperplane is fixed, we must consider the distribution of the vertex instead. A non-boundary vertex of $\Xi(\mathcal{S})$ is defined by the intersection of h_1 hyperplanes associated with one simplex, h_2 hyperplanes associated with another and so on, where $\sum_i h_i = d$. If v lies in a small region of volume $\mathrm{d}\sigma$ with centre p and diameter δ, then *all* of these hyperplanes must pass through that region. There are two possible cases for the set of h hyperplanes associated with a particular simplex s:

Case 1: The hyperplanes are all walls supporting the simplex. Their intersection is a $(d-h)$-dimensional "wall" Z standing on the affine span of $d-h$ vertices of s. Theorem 10, with $k = d-h$ and $k' = 0$, tells us that the probability that Z is δ-close to p is at most $K\delta^h\gamma^{d-1}$.

Case 2: The hyperplanes include the affine span of the simplex itself. Then their intersection is simply the affine span of a $(d-h)$-face of the simplex and Theorem 7, with $k = d-h$, tells us that the probability that Z is δ-close to p is at most $K\delta^h\gamma^{d-h+1}$.

So the probability that all the hyperplanes pass through $\mathrm{d}\sigma$ is at most

$$\begin{aligned}\prod_i K_i\delta^{h_i}\max\{\gamma,\gamma^d\} &\leq K'\delta^d\max\{\gamma,\gamma^{d^2}\}\\ &\leq K''\mathrm{d}\sigma\max\{\gamma,\gamma^{d^2}\}\end{aligned}$$

(The last step assumes that the small region is "round", i.e. $\mathrm{d}\sigma = \Theta(\delta^d)$.) In other words, the vertex v follows a $K''\max\{\gamma,\gamma^{d^2}\}$-smooth distribution. The portion of the domain $\mathcal{C}$ within distance ε of the hyperplane H has volume at most $\varepsilon V_{d-1}(1)$, so the probability that v is ε-close to H is at most $K''\varepsilon\max\{\gamma,\gamma^{d^2}\}V_{d-1}(1)$.

If the vertex lies on the boundary, we must modify our analysis only slightly. Constrain ε to be less than D_{in}, and assume b hyperplanes from the boundary contain v. If $h_1, h_2, \ldots$ hyperplanes associated with simplices also contain v as before, then it must be the case that $\sum_i h_i = d-b$. The b hyperplanes on the boundary intersect in a $(d-b)$-flat B. Carry out the previous analysis assuming that the differential region is a subset of B and not of the full-dimensional space. Then $\mathrm{d}\sigma = \Theta(\delta^{d-b})$ and we obtain the result that v is $K'''\max\{\gamma,\gamma^{d(d-b)}\}$-smoothly distributed on B. The boundary is fixed, so B and H are at a constant angle. If this angle is zero (B and H are parallel), then B is D_{in}-distant, and hence ε-distant, from H. Else, the $(d-b)$-measure of the region of B within $\mathcal{C}$ ε-close to H is at most $C\varepsilon V_{d-b-1}(1)$ for some constant C. The probability that v lies in this region is at most $K'''C\varepsilon\max\{\gamma,\gamma^{d(d-b)}\}V_{d-b-1}(1)$.

This concludes the proof of Theorem 4.

4. Conclusion

The following extensions of the presented analysis naturally suggest themselves and are left for future work.

- **Translational motion planning.** The free configuration space for translational motion planning of a polyhedral robot among polyhedral obstacles is the complement of the union of Minkowski sums of the obstacles and the antipode of the robot [35]. Independently perturbed simplices cannot model this setting since the connectivity of the Minkowski sums is not preserved. Extending our results to translational motion planning necessitates relieving the reliance of the current analysis on complete independence of the perturbations. Preliminary derivations suggest that the limited amount of independence present in the free configuration space of translational motion planning is sufficient to obtain a polynomial bound on the number of milestones.
- **General motion planning and curved $\mathcal{C}$-space obstacles.** General motion planning, such as holonomic or articulated motion with translations and rotations, gives rise to configuration spaces with curved $\mathcal{C}$-space obstacles, generally represented as semi-algebraic sets. In order to do smoothed analysis in this setting, a convincing perturbation model for semi-algebraic sets needs to be defined. Building on this definition, a polynomial number of random samples has to be shown to yield an accurate roadmap.
- **Connection to previous theoretical models.** It is reasonable to conjecture that smoothly perturbed free configuration spaces are $(\varepsilon, \alpha, \beta)$-expansive, for appropriate values of ε, α, and β [22]. This would imply that results previously obtained for expansive configuration spaces carry over to the smoothed setting.

Acknowledgements

We are grateful to Dan Halperin for crucial early discussions and to Jean-Claude Latombe for his support throughout this project. We also acknowledge the comments of the anonymous referees: we particularly appreciate the careful reading that revealed a flaw in a revised proof, and suggestions that helped us improve the exposition.

References

[1] P. K. Agarwal, M. J. Katz, and M. Sharir. Computing depth orders for fat objects and related problems. *Computational Geometry Theory and Applications*, 5:187–206, 1995.

[2] B. Aronov and M. Sharir. Castles in the air revisited. *Discrete and Computational Geometry*, 12:119–150, 1994.

[3] D. Arthur and S. Vassilvitskii. Worst-case and smoothed analyses of the ICP algorithm, with an application to the k-means method. In *Proc. 47th IEEE Symposium on Foundations of Computer Science*, 2006.

[4] C. Banderier, R. Beier, and K. Mehlhorn. Smoothed analysis of three combinatorial problems. In *Proc. 28th Symposium on Mathematical Foundations of Computer Science*, pages 198–207, 2003.

[5] J. Barraquand, L. E. Kavraki, J.-C. Latombe, T.-Y. Li, R. Motwani, and P. Raghavan. A random sampling scheme for path planning. *International Journal of Robotics Research*, 16:759–774, 1997.

[6] R.-P. Berretty, M. Overmars, and A. F. van der Stappen. Dynamic motion planning in low obstacle density environments. In *Proc. Workshop Algorithms Data Struct.*, volume 1272 of *Lecture Notes in Computer Science*, pages 3–16. Springer-Verlag, 1997.
[7] A. Blum and J. Dunagan. Smoothed analysis of the perceptron algorithm for linear programming. In *Proc. 13th ACM-SIAM Symposium on Discrete Algorithms*, pages 905–914, 2002.
[8] D. E. Cardoze and L. J. Schulman. Pattern matching for spatial point sets. In *Proc. 39th IEEE Symposium on Foundations of Computer Science*, pages 156–165, 1998.
[9] B. Chazelle, H. Edelsbrunner, L. J. Guibas, and M. Sharir. A singly-exponential stratification scheme for real semi-algebraic varieties and its applications. *Theoretical Computer Science*, 84:77–105, 1991.
[10] K. L. Clarkson. Nearest neighbor queries in metric spaces. *Discrete and Computational Geometry*, 22:63–93, 1999.
[11] V. Damerow, F. M. auf der Heide, H. Räcke, C. Scheideler, and C. Sohler. Smoothed motion complexity. In *Proc. 11th European Symposium on Algorithms*, pages 161–171, 2003.
[12] V. Damerow and C. Sohler. Extreme points under random noise. In *Proc. 12th European Symposium on Algorithms*, pages 264–274, 2004.
[13] M. de Berg. Linear size binary space partitions for fat objects. In *Proc. 3rd Annu. European Sympos. Algorithms*, volume 979 of *Lecture Notes in Computer Science*, pages 252–263. Springer-Verlag, 1995.
[14] M. de Berg. Linear size binary space partitions for uncluttered scenes. *Algorithmica*, 28:353–366, 2000.
[15] M. de Berg, D. Halperin, M. Overmars, and M. van Kreveld. Sparse arrangements and the number of views of polyhedral scenes. *International Journal of Computational Geometry and Applications*, 7:175–195, 1997.
[16] M. de Berg, M. J. Katz, A. F. van der Stappen, and J. Vleugels. Realistic input models for geometric algorithms. In *Proc. 13th Symposium on Computational Geometry*, pages 294–303, 1997.
[17] A. Deshpande and D. A. Spielman. Improved smoothed analysis of the shadow vertex simplex method. In *Proc. 46th IEEE Symposium on Foundations of Computer Science*, pages 349–356, 2005.
[18] J. Erickson. Nice point sets can have nasty delaunay triangulations. *Discrete and Computational Geometry*, 30:109–132, 2003.
[19] M. Gavrilov, P. Indyk, R. Motwani, and S. Venkatasubramanian:. Geometric pattern matching: A performance study. In *Proc. 15th Symposium on Computational Geometry*, pages 79–85, 1999.
[20] R. Geraerts and M. Overmars. A comparative study of probabilistic roadmap planners. In *Proc. 5th Workshop on the Algorithmic Foundations of Robotics*, pages 43–59, 2002.
[21] D. Hsu, J.-C. Latombe, and H. Kurniawati. On the probabilistic foundations of probabilistic roadmap planning. *International Journal of Robotics Research.* to appear.
[22] D. Hsu, J.-C. Latombe, and R. Motwani. Path planning in expansive configuration spaces. *International Journal of Computational Geometry and Applications*, 9:495–512, 1999.
[23] D. Hsu, J.-C. Latombe, R. Motwani, and L. E. Kavraki. Capturing the connectivity of high-dimensional geometric spaces by parallelizable random sampling techniques. In P. Pardalos and S. Rajasekaran, editors, *Advances in Randomized Parallel Computing*, pages 159–182. Kluwer Academic Publishers, Boston, MA, 1999.
[24] L. E. Kavraki, M. N. Kolountzakis, and J.-C. Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Transactions on Robotics and Automation*, 14:166–171, 1998.
[25] L. E. Kavraki and J.-C. Latombe. Probabilistic roadmaps for robot path planning. In K. Gupta and A. del Pobil, editors, *Practical Motion Planning in Robotics: Current Approaches and Future Directions*, pages 33–53. John Wiley, 1998.
[26] L. E. Kavraki, J.-C. Latombe, R. Motwani, and P. Raghavan. Randomized query processing in robot path planning. In *Proc. 27th ACM Symposium on Theory of Computing*, pages 353–362, 1995.
[27] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 11:566–580, 1996.
[28] V. Koltun. Almost tight upper bounds for vertical decompositions in four dimensions. *Journal of the ACM*, 51:699–730, 2004.
[29] J.-C. Latombe. *Robot Motion Planning.* Kluwer Academic Publishers, Boston, 1991.
[30] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. E. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital Michelangelo project: 3D scanning of

large statues. In *SIGGRAPH*, pages 131–144, 2000.

[31] R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge University Press, New York, NY, 1995.

[32] M. H. Overmars and A. F. van der Stappen. Range searching and point location among fat objects. *Journal of Algorithms*, 21:629–656, 1996.

[33] G. Pisier. *The volume of convex bodies and Banach space geometry.* Cambridge Tracts in Mathematics, 94, Cambridge University Press., 1989.

[34] O. Schwarzkopf and J. Vleugels. Range searching in low-density environments. *Information Processing Letters*, 60:121–127, 1996.

[35] M. Sharir. Algorithmic motion planning. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 40, pages 733–754. CRC Press LLC, Boca Raton, FL, 1997.

[36] G. Song, S. L. Thomas, and N. M. Amato. A general framework for PRM motion planning. In *Proc. International Conference on Robotics and Automation*, pages 4445–4450, 2003.

[37] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms. In *Proc. of the International Congress of Mathematicians 2002.*

[38] D. A. Spielman and S.-H. Teng. Smoothed analysis of termination of linear programming algorithms. *Mathematical Programming, Series B*, 97, 2003.

[39] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51:385–463, 2004.

[40] A. F. van der Stappen and M. H. Overmars. Motion planning amidst fat obstacles. In *Proc. 10th Symposium on Computational Geometry*, pages 31–40, 1994.