Chapter 19

The Johnson-Lindenstrauss Lemma

By Sariel Har-Peled, December 30, 2015[①]

> Dixon was alive again. Consciousness was upon him before he could get out of the way; not for him the slow, gracious wandering from the halls of sleep, but a summary, forcible ejection. He lay sprawled, too wicked to move, spewed up like a broken spider-crab on the tarry shingle of the morning. The light did him harm, but not as much as looking at things did; he resolved, having done it once, never to move his eyeballs again. A dusty thudding in his head made the scene before him beat like a pulse. His mouth had been used as a latrine by some small creature of the night, and then as its mausoleum. During the night, too, he'd somehow been on a cross-country run and then been expertly beaten up by secret police. He felt bad.
>
> – Lucky Jim, Kingsley Amis.

In this chapter, we will prove that given a set P of n points in $\mathbb{R}^d$, one can reduce the dimension of the points to $k = O(\varepsilon^{-2} \log n)$ such that distances are $1 \pm \varepsilon$ preserved. Surprisingly, this reduction is done by randomly picking a subspace of k dimensions and projecting the points into this random subspace. One way of thinking about this result is that we are "compressing" the input of size nd (i.e., n points with d coordinates) into size $O(n\varepsilon^{-2} \log n)$, while (approximately) preserving distances.

19.1. The Brunn-Minkowski inequality

For a set $A \subseteq \mathbb{R}^d$, an a point $\mathsf{p} \in \mathbb{R}^d$, let $A + \mathsf{p}$ denote the translation of A by p. Formally, $A + \mathsf{p} = \left\{ \mathsf{q} + \mathsf{p} \;\middle|\; \mathsf{q} \in A \right\}$.

Definition 19.1.1. For two sets A and B in $\mathbb{R}^n$, let $A + B$ denote the *Minkowski sum* of A and B. Formally,

$$A + B = \left\{ a + b \;\middle|\; a \in A, b \in B \right\} = \bigcup_{\mathsf{p} \in A} (\mathsf{p} + B).$$

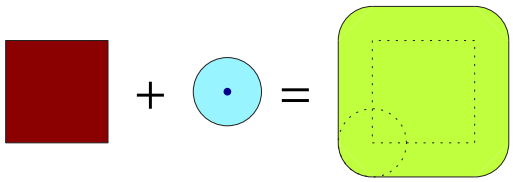

Remark 19.1.2. It is easy to verify that if A' and B' are translated copies of A and B (that is, $A' = A + \mathsf{p}$ and $B = B + \mathsf{q}$, for some points $\mathsf{p}, \mathsf{q} \in \mathbb{R}^d$), respectively, then $A' + B'$ is a translated copy of $A + B$. In particular, since volume is preserved under translation, we have that $\mathrm{vol}\Big(A' + B'\Big) = \mathrm{vol}\Big((A + B) + \mathsf{p} + \mathsf{q}\Big) = \mathrm{vol}\Big(A + B\Big)$, where $\mathrm{vol}(X)$ is the *volume* (i.e., measure) of the set X.

[①]This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nc/3.0/ or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

Our purpose here is to prove the following theorem.

Theorem 19.1.3 (Brunn-Minkowski inequality). *Let A and B be two non-empty compact sets in $\mathbb{R}^n$. Then*

$$\operatorname{vol}(A+B)^{1/n} \geq \operatorname{vol}(A)^{1/n} + \operatorname{vol}(B)^{1/n}.$$

Definition 19.1.4. A set $A \subseteq \mathbb{R}^n$ is a *brick set* if it is the union of finitely many (close) axis parallel boxes with disjoint interiors.

It is intuitively clear, by limit arguments, that proving Theorem 19.1.3 for brick sets will imply it for the general case.

Lemma 19.1.5 (Brunn-Minkowski inequality for Brick Sets). *Let A and B be two non-empty brick sets in $\mathbb{R}^n$. Then*

$$\Big(\operatorname{vol}(A+B)\Big)^{1/n} \geq \operatorname{vol}(A)^{1/n} + \operatorname{vol}(B)^{1/n}.$$

Proof: By induction on the number k of bricks in A and B. If $k = 2$ then A and B are just bricks, with dimensions $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$, respectively. In this case, the dimensions of $A + B$ are $a_1 + b_1, \ldots, a_n + b_n$, as can be easily verified. Thus, we need to prove that $(\prod_{i=1}^n a_i)^{1/n} + (\prod_{i=1}^n b_i)^{1/n} \leq \Big(\prod_{i=1}^n (a_i + b_i)\Big)^{1/n}$.Dividing the left side by the right side, we have

$$\left(\prod_{i=1}^n \frac{a_i}{a_i+b_i}\right)^{1/n} + \left(\prod_{i=1}^n \frac{b_i}{a_i+b_i}\right)^{1/n} \leq \frac{1}{n}\sum_{i=1}^n \frac{a_i}{a_i+b_i} + \frac{1}{n}\sum_{i=1}^n \frac{b_i}{a_i+b_i} = 1,$$

by the generalized arithmetic-geometric mean inequality[②], and the claim follows for this case.

Now let $k > 2$ and suppose that the Brunn-Minkowski inequality holds for any pair of brick sets with fewer than k bricks (together). Let A and B be a pair of sets having k bricks together, the A has at least two (disjoint) bricks. However, this implies that there is an axis parallel hyperplane h that separates the interior of one brick of A from the interior of another brick of A (the hyperplane h might intersect other bricks of A). Assume that h is the hyperplane $x_1 = 0$ (this can be achieved by translation and renaming of coordinates).

Let $\overline{A^+} = A \cap h^+$ and $\overline{A^-} = A \cap h^-$, where h^+ and h^- are the two open half spaces induced by h. Let A^+ and A^- be the closure of $\overline{A^+}$ and $\overline{A^-}$, respectively. Clearly, A^+ and A^- are both brick sets with (at least) one fewer brick than A.

Next, observe that the claim is translation invariant (see Remark 19.1.2), and as such, let us translate B so that its volume is split by h in the same ratio A's volume is being split. Denote the two parts of B by B^+ and B^-, respectively. Let $\rho = \operatorname{vol}(A^+)/\operatorname{vol}(A) = \operatorname{vol}(B^+)/\operatorname{vol}(B)$ (if $\operatorname{vol}(A) = 0$ or $\operatorname{vol}(B) = 0$ the claim trivially holds).

Observe, that $A^+ + B^+ \subseteq A + B$, and it lies on one side of h (since $h \equiv (x_1 = 0)$), and similarly $A^- + B^- \subseteq A + B$ and it lies on the other side of h. Thus, by induction and since $A^+ + B^+$ and $A^- + B^-$ are interior disjoint, we have

$$\begin{aligned}\operatorname{vol}(A+B) &\geq \operatorname{vol}(A^+ + B^+) + \operatorname{vol}(A^- + B^-) \\ &\geq \Big(\operatorname{vol}(A^+)^{1/n} + \operatorname{vol}(B^+)^{1/n}\Big)^n + \Big(\operatorname{vol}(A^-)^{1/n} + \operatorname{vol}(B^-)^{1/n}\Big)^n\end{aligned}$$

[②]Here is a proof of the generalized form: Let $x_1, \ldots, x_n$ be n positive real numbers. Consider the quantity $R = x_1 x_2 \cdots x_n$. If we fix the sum of the n numbers to be equal to α, then R is maximized when all the x_is are equal. Thus, $\sqrt[n]{x_1 x_2 \cdots x_n} \leq \sqrt[n]{(\alpha/n)^n} = \alpha/n = (x_1 + \cdots + x_n)/n$.

$$
\begin{aligned}
&= \Big[\rho^{1/n}\mathrm{vol}(A)^{1/n} + \rho^{1/n}\mathrm{vol}(B)^{1/n}\Big]^n \\
&\qquad \Big[(1-\rho)^{1/n}\mathrm{vol}(A)^{1/n} + (1-\rho)^{1/n}\mathrm{vol}(B)^{1/n}\Big]^n \\
&= (\rho + (1-\rho))\Big[\mathrm{vol}(A)^{1/n} + \mathrm{vol}(B)^{1/n}\Big]^n \\
&= \Big[\mathrm{vol}(A)^{1/n} + \mathrm{vol}(B)^{1/n}\Big]^n,
\end{aligned}
$$

establishing the claim. ■

Proof of ***Theorem 19.1.3***: Let $A_1 \subseteq A_2 \subseteq \cdots \subseteq A_i \subseteq \cdots$ be a sequence of finite brick sets, such that $\bigcup_i A_i = A$, and similarly let $B_1 \subseteq B_2 \subseteq \cdots \subseteq B_i \subseteq \cdots$ be a sequence of finite brick sets, such that $\bigcup_i B_i = B$. By the definition of volume[3],we have that $\lim_{i\to\infty} \mathrm{vol}(A_i) = \mathrm{vol}(A)$ and $\lim_{i\to\infty} \mathrm{vol}(B_i) = \mathrm{vol}(B)$.

We claim that $\lim_{i\to\infty} \mathrm{vol}(A_i + B_i) = \mathrm{vol}(A + B)$. Indeed, consider any point $z \in A + B$, and let $u \in A$ and $v \in B$ be such that $u + v = z$. By definition, there exists an i, such that for all $j > i$ we have $u \in A_j$, $v \in B_j$, and as such $z \in A_j + B_j$. Thus, $A + B \subseteq \cup_j (A_j + B_j)$ and $\cup_j (A_j + B_j) \subseteq \cup_j (A + B) \subseteq A + B$; namely, $\cup_j (A_j + B_j) = A + B$.

Furthermore, for any $i > 0$, since A_i and B_i are brick sets, we have

$$
\mathrm{vol}(A_i + B_i)^{1/n} \geq \mathrm{vol}(A_i)^{1/n} + \mathrm{vol}(B_i)^{1/n},
$$

by Lemma 19.1.5. Thus,

$$
\begin{aligned}
\mathrm{vol}\Big(A + B\Big)^{1/n} &= \lim_{i\to\infty} \mathrm{vol}(A_i + B_i)^{1/n} \geq \lim_{i\to\infty} \Big(\mathrm{vol}(A_i)^{1/n} + \mathrm{vol}(B_i)^{1/n}\Big) \\
&= \mathrm{vol}(A)^{1/n} + \mathrm{vol}(B)^{1/n}.
\end{aligned}
$$

■

Theorem 19.1.6 (Brunn-Minkowski for slice volumes.). *Let $\mathcal{P}$ be a convex set in $\mathbb{R}^{n+1}$, and let $A = \mathcal{P} \cap (x_1 = a)$, $B = \mathcal{P} \cap (x_1 = b)$ and $C = \mathcal{P} \cap (x_1 = c)$ be three slices of A, for $a < b < c$. We have* $\mathrm{vol}(B) \geq \min(\mathrm{vol}(A), \mathrm{vol}(C))$. *Specifically, consider the function*

$$
v(t) = \left(\mathrm{vol}\Big(\mathcal{P} \cap (x_1 = t)\Big)\right)^{1/n},
$$

and let $\mathcal{I} = [t_{min}, t_{max}]$ be the interval where the hyperplane $x_1 = t$ intersects $\mathcal{P}$. Then, $v(t)$ is concave on $\mathcal{I}$.

Proof: If a or c are outside $\mathcal{I}$, then $\mathrm{vol}(A) = 0$ or $\mathrm{vol}(C) = 0$, respectively, and then the claim trivially holds.

Otherwise, let $\alpha = (b - a)/(c - a)$. We have that $b = (1 - \alpha) \cdot a + \alpha \cdot c$, and by the convexity of $\mathcal{P}$, we have $(1 - \alpha)A + \alpha C \subseteq B$. Thus, by Theorem 19.1.3 we have

$$
\begin{aligned}
v(b) = \mathrm{vol}(B)^{1/n} &\geq \mathrm{vol}((1-\alpha)A + \alpha C)^{1/n} \geq \mathrm{vol}((1-\alpha)A)^{1/n} + \mathrm{vol}(\alpha C)^{1/n} \\
&= ((1-\alpha)^n \mathrm{vol}(A))^{1/n} + (\alpha^n \mathrm{vol}(C))^{1/n} \\
&= (1-\alpha) \cdot \mathrm{vol}(A)^{1/n} + \alpha \cdot \mathrm{vol}(C)^{1/n} \\
&= (1-\alpha)v(a) + \alpha v(c).
\end{aligned}
$$

Namely, $v(\cdot)$ is concave on $\mathcal{I}$, and in particular $v(b) \geq \min(v(a), v(c))$, which in turn implies that $\mathrm{vol}(B) = v(b)^n \geq (\min(v(a), v(c)))^n = \min(\mathrm{vol}(A), \mathrm{vol}(C))$, as claimed. ■

[3] This is the standard definition in measure theory of volume. The reader unfamiliar with this fanfare can either consult a standard text on the topic, or take it for granted as this is intuitively clear.

Corollary 19.1.7. *For A and B compact sets in $\mathbb{R}^n$, the following holds* $\mathrm{vol}((A+B)/2) \geq \sqrt{\mathrm{vol}(A)\mathrm{vol}(B)}$.

Proof: We have that

$$\begin{aligned}\mathrm{vol}\Big((A+B)/2\Big)^{1/n} &= \mathrm{vol}\Big(A/2+B/2\Big)^{1/n} \geq \mathrm{vol}\Big(A/2\Big)^{1/n} + \mathrm{vol}\Big(B/2\Big)^{1/n} = \Big(\mathrm{vol}(A)^{1/n} + \mathrm{vol}(B)^{1/n}\Big)/2 \\ &\geq \sqrt{\mathrm{vol}(A)^{1/n}\mathrm{vol}(B)^{1/n}}\end{aligned}$$

by Theorem 19.1.3, and since $(a+b)/2 \geq \sqrt{ab}$ for any $a, b \geq 0$. The claim now follows by raising this inequality to the power n. ■

19.1.1. The Isoperimetric Inequality

The following is not used anywhere else and is provided because of its mathematical elegance. The skip-able reader can thus employ their special gift and move on to Section 19.2.

The *isoperimetric inequality* states that among all convex bodies of a fixed surface area, the ball has the largest volume (in particular, the unit circle is the largest area planar region with perimeter 2π). This problem can be traced back to antiquity, in particular Zenodorus (200–140 BC) wrote a monograph (which was lost) that seemed to have proved the claim in the plane for some special cases. The first formal proof for the planar case was done by Steiner in 1841. Interestingly, the more general claim is an easy consequence of the Brunn-Minkowski inequality.

Let K be a convex body in $\mathbb{R}^n$ and $\mathbf{b}$ be the n dimensional ball of radius one centered at the origin. Let $\mathsf{S}(X)$ denote the surface area of a compact set $X \subseteq \mathbb{R}^n$. The *isoperimetric inequality* states that

$$\left(\frac{\mathrm{vol}(K)}{\mathrm{vol}(\mathbf{b})}\right)^{1/n} \leq \left(\frac{\mathsf{S}(K)}{\mathsf{S}(\mathbf{b})}\right)^{1/(n-1)}. \tag{19.1}$$

Namely, the left side is the radius of a ball having the same volume as K, and the right side is the radius of a sphere having the same surface area as K. In particular, if we scale K so that its surface area is the same as $\mathbf{b}$, then the above inequality implies that $\mathrm{vol}(K) \leq \mathrm{vol}(\mathbf{b})$.

To prove Eq. (19.1), observe that $\mathrm{vol}(\mathbf{b}) = \mathsf{S}(\mathbf{b})/n$[④]. Also, observe that $K + \varepsilon\mathbf{b}$ is the body K together with a small "atmosphere" around it of thickness ε. In particular, the volume of this "atmosphere" is (roughly) $\varepsilon\,\mathsf{S}(K)$ (in fact, Minkowski defined the surface area of a convex body to be the limit stated next).

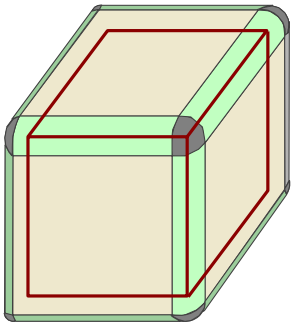

Formally, we have

$$\begin{aligned}\mathsf{S}(K) &= \lim_{\varepsilon\to 0+} \frac{\mathrm{vol}(K+\varepsilon\mathbf{b}) - \mathrm{vol}(K)}{\varepsilon} \\ &\geq \lim_{\varepsilon\to 0+} \frac{\Big(\mathrm{vol}(K)^{1/n} + \mathrm{vol}(\varepsilon\mathbf{b})^{1/n}\Big)^n - \mathrm{vol}(K)}{\varepsilon},\end{aligned}$$

by the Brunn-Minkowski inequality. Now $\mathrm{vol}(\varepsilon\mathbf{b})^{1/n} = \varepsilon\mathrm{vol}(\mathbf{b})^{1/n}$, and as such

$$\begin{aligned}\mathsf{S}(K) &\geq \lim_{\varepsilon\to 0+} \frac{\mathrm{vol}(K) + \binom{n}{1}\varepsilon\mathrm{vol}(K)^{(n-1)/n}\mathrm{vol}(\mathbf{b})^{1/n} + \binom{n}{2}\varepsilon^2\langle\cdots\rangle + \cdots + \varepsilon^n\mathrm{vol}(\mathbf{b}) - \mathrm{vol}(K)}{\varepsilon} \\ &= \lim_{\varepsilon\to 0+} \frac{n\varepsilon\mathrm{vol}(K)^{(n-1)/n}\mathrm{vol}(\mathbf{b})^{1/n}}{\varepsilon} = n\mathrm{vol}(K)^{(n-1)/n}\mathrm{vol}(\mathbf{b})^{1/n}.\end{aligned}$$

[④]Indeed, $\mathrm{vol}(\mathbf{b}) = \int_{r=0}^{1} \mathsf{S}(\mathbf{b}) r^{n-1} dr = \mathsf{S}(\mathbf{b})/n$.

Dividing both sides by $\mathsf{S}(\mathbf{b}) = n\mathrm{vol}(\mathbf{b})$, we have

$$\frac{\mathsf{S}(K)}{\mathsf{S}(\mathbf{b})} \geq \frac{\mathrm{vol}(K)^{(n-1)/n}}{\mathrm{vol}(\mathbf{b})^{(n-1)/n}} \qquad \Longrightarrow \qquad \left(\frac{\mathsf{S}(K)}{\mathsf{S}(\mathbf{b})}\right)^{1/(n-1)} \geq \left(\frac{\mathrm{vol}(K)}{\mathrm{vol}(\mathbf{b})}\right)^{1/n},$$

establishing the isoperimetric inequality.

19.2. Measure Concentration on the Sphere

Let $\mathbb{S}^{(n-1)}$ be the unit sphere in $\mathbb{R}^n$. We assume there is a uniform probability measure defined over $\mathbb{S}^{(n-1)}$, such that its total measure is 1. Surprisingly, most of the mass of this measure is near the equator. Indeed, consider an arbitrary equator π on $\mathbb{S}^{(n-1)}$ (that it, it is the intersection of the sphere with a hyperplane passing through the center of ball inducing the sphere). Next, consider all the points that are in distance $\approx \ell(n) = c/n^{1/3}$ from π. The question we are interested in is what fraction of the sphere is covered by this strip T (depicted on the right).

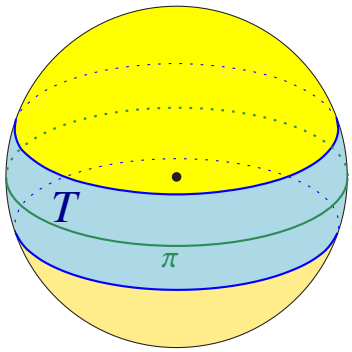

Notice, that as the dimension increases the width $\ell(n)$ of this strip decreases. But surprisingly, despite its width becoming smaller, as the dimension increases, this strip contains a larger and larger fraction of the sphere. In particular, the total fraction of the sphere not covered by this (shrinking!) strip converges to zero.

Furthermore, counter intuitively, this is true for *any* equator. We are going to show that even a stronger result holds: The mass of the sphere is concentrated close to the boundary of any set $A \subseteq \mathbb{S}^{(n-1)}$ such that $\mathbf{Pr}[A] = 1/2$.

Before proving this somewhat surprising theorem, we will first try to get an intuition about the behavior of the hypersphere in high dimensions.

19.2.1. The strange and curious life of the hypersphere

Consider the ball of radius r in $\mathbb{R}^n$ denoted by $r\,\mathbf{b}^n$, where $\mathbf{b}^n$ is the unit radius ball centered at the origin. Clearly, $\mathrm{vol}(r\,\mathbf{b}^n) = r^n\mathrm{vol}(\mathbf{b}^n)$. Now, even if r is very close to 1, the quantity r^n might be very close to zero if n is sufficiently large. Indeed, if $r = 1-\delta$, then $r^n = (1-\delta)^n \leq \exp(-\delta n)$, which is very small if $\delta \gg 1/n$. (Here, we used the fact that $1 - x \leq e^x$, for $x \geq 0$.) Namely, for the ball in high dimensions, its mass is concentrated in a very thin shell close to its surface.

The volume of a ball and the surface area of hypersphere. Let $\mathrm{vol}(r\mathbf{b}^n)$ denote the volume of the ball of radius r in $\mathbb{R}^n$, and $\mathrm{Area}\left(r\mathbb{S}^{(n-1)}\right)$ denote the surface area of its bounding sphere (i.e., the surface area of $r\mathbb{S}^{(n-1)}$). It is known that

$$\mathrm{vol}(r\mathbf{b}^n) = \frac{\pi^{n/2}r^n}{\Gamma(n/2+1)} \quad \text{and} \quad \mathrm{Area}\left(r\mathbb{S}^{(n-1)}\right) = \frac{2\pi^{n/2}r^{n-1}}{\Gamma(n/2)},$$

where the gamma function, $\Gamma(\cdot)$, is an extension of the factorial function. Specifically, if n is even then $\Gamma(n/2+1) = (n/2)!$, and for n odd $\Gamma(n/2+1) = \sqrt{\pi}(n!!)/2^{(n+1)/2}$, where $n!! = 1\cdot 3\cdot 5\cdots n$ is the *double factorial*. The most surprising implication of these two formulas is that, as n increases, the volume of the unit ball first increases (till dimension 5 in fact) and then starts decreasing to zero.

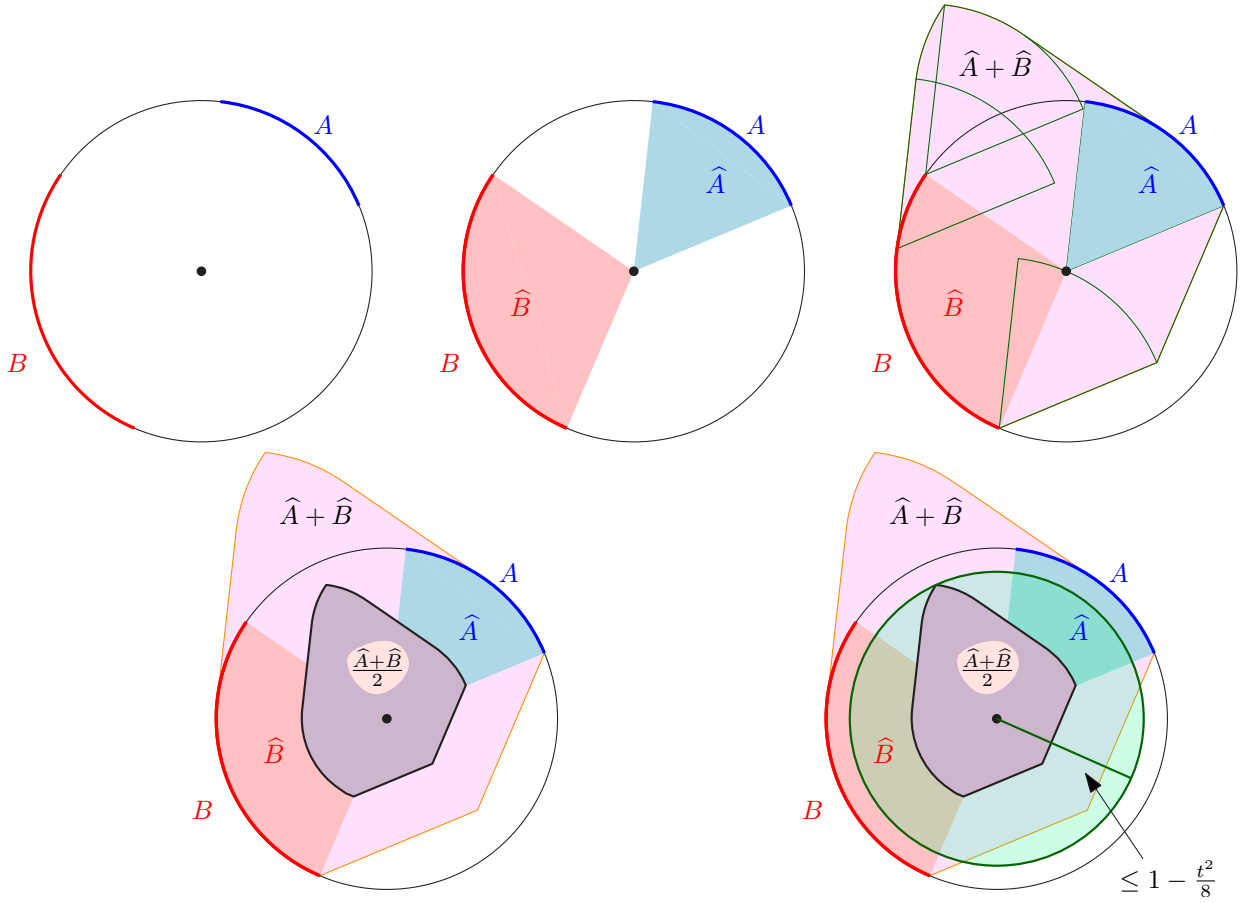

Figure 19.1: Illustration of the proof of Theorem 19.2.1.

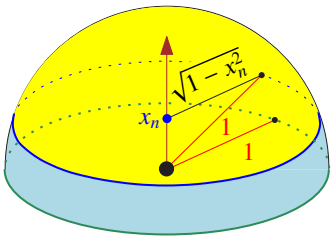

Similarly, the surface area of the unit sphere $\mathbb{S}^{(n-1)}$ in $\mathbb{R}^n$ tends to zero as the dimension increases. To see this, compute the volume of the unit ball using an integral of its slice volume, when it is being sliced by a hyperplanes perpendicular to the nth coordinate.

We have, see figure on the right, that

$$\operatorname{vol}\left(\mathbf{b}^n\right) = \int_{x_n=-1}^{1} \operatorname{vol}\left(\sqrt{1-x_n^2}\,\mathbf{b}^{n-1}\right) dx_n = \operatorname{vol}\left(\mathbf{b}^{n-1}\right) \int_{x_n=-1}^{1} \left(1-x_n^2\right)^{(n-1)/2} dx_n,$$

Now, the integral on the right side tends to zero as n increases. In fact, for n very large, the term $\left(1-x_n^2\right)^{(n-1)/2}$ is very close to 0 everywhere except for a small interval around 0. This implies that the main contribution of the volume of the ball happens when we consider slices of the ball by hyperplanes of the form $x_n = \delta$, where δ is small.

If one has to visualize how such a ball in high dimensions looks like, it might be best to think about it as a star-like creature: It has very little mass close to the tips of any set of orthogonal directions we pick, and most of its mass somehow lies on hyperplanes close to its center.[5]

[5] In short, it looks like a Boojum [Car76].

19.2.2. Measure Concentration on the Sphere

Theorem 19.2.1 (Measure concentration on the sphere.). *Let $A \subseteq \mathbb{S}^{(n-1)}$ be a measurable set with $\mathbf{Pr}[A] \geq 1/2$, and let A_t denote the set of points of $\mathbb{S}^{(n-1)}$ in distance at most t from A, where $t \leq 2$. Then $1 - \mathbf{Pr}[A_t] \leq 2\exp\left(-nt^2/2\right)$.*

Proof: We will prove a slightly weaker bound, with $-nt^2/4$ in the exponent. Let $\widehat{A} = T(A)$, where

$$T(X) = \left\{\alpha x \;\middle|\; x \in X, \alpha \in [0,1]\right\} \subseteq \mathbf{b}^n,$$

and $\mathbf{b}^n$ is the unit ball in $\mathbb{R}^n$. We have that $\mathbf{Pr}[A] = \mu\left(\widehat{A}\right)$, where $\mu\left(\widehat{A}\right) = \mathrm{vol}\left(\widehat{A}\right)/\mathrm{vol}(\mathbf{b}^n)$[6].

Let $B = \mathbb{S}^{(n-1)} \setminus A_t$ and $\widehat{B} = T(B)$, see Figure 19.1. We have that $\|a - b\| \geq t$ for all $a \in A$ and $b \in B$. By Lemma 19.2.2 below, the set $\left(\widehat{A} + \widehat{B}\right)/2$ is contained in the ball $r\mathbf{b}^n$ centered at the origin, where $r = 1 - t^2/8$. Observe that $\mu(r\mathbf{b}^n) = \mathrm{vol}(r\mathbf{b}^n)/\mathrm{vol}(\mathbf{b}^n) = r^n = \left(1 - t^2/8\right)^n$. As such, applying the Brunn-Minkowski inequality in the form of Corollary 19.1.7, we have

$$\left(1 - \frac{t^2}{8}\right)^n = \mu\left(r\mathbf{b}^n\right) \geq \mu\left(\frac{\widehat{A} + \widehat{B}}{2}\right) \geq \sqrt{\mu\left(\widehat{A}\right)\mu\left(\widehat{B}\right)} = \sqrt{\mathbf{Pr}[A]\,\mathbf{Pr}[B]} \geq \sqrt{\mathbf{Pr}[B]\,/2}\,.$$

Thus, $\mathbf{Pr}[B] \leq 2(1 - t^2/8)^{2n} \leq 2\exp(-2nt^2/8)$, since $1 - x \leq \exp(-x)$, for $x \geq 0$. ■

Lemma 19.2.2. *For any $\widehat{a} \in \widehat{A}$ and $\widehat{b} \in \widehat{B}$, we have* $\left\|\frac{\widehat{a} + \widehat{b}}{2}\right\| \leq 1 - \frac{t^2}{8}$.

Proof: Let $\widehat{a} = \alpha a$ and $\widehat{b} = \beta b$, where $a \in A$ and $b \in B$. We have

$$\|u\| = \left\|\frac{a+b}{2}\right\| = \sqrt{1^2 - \left\|\frac{a-b}{2}\right\|^2} \leq \sqrt{1 - \frac{t^2}{4}} \leq 1 - \frac{t^2}{8}, \tag{19.2}$$

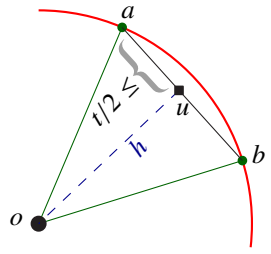

since $\|a - b\| \geq t$. As for $\widehat{a}$ and $\widehat{b}$, assume that $\alpha \leq \beta$, and observe that the quantity $\left\|\widehat{a} + \widehat{b}\right\|$ is maximized when $\beta = 1$. As such, by the triangle inequality, we have

$$\begin{aligned}\left\|\frac{\widehat{a} + \widehat{b}}{2}\right\| &= \left\|\frac{\alpha a + b}{2}\right\| \leq \left\|\frac{\alpha(a+b)}{2}\right\| + \left\|(1-\alpha)\frac{b}{2}\right\| \\ &\leq \alpha\left(1 - \frac{t^2}{8}\right) + (1-\alpha)\frac{1}{2} = \tau,\end{aligned}$$

by Eq. (19.2) and since $\|b\| = 1$. Now, τ is a convex combination of the two numbers $1/2$ and $1 - t^2/8$. In particular, we conclude that $\tau \leq \max(1/2, 1 - t^2/8) \leq 1 - t^2/8$, since $t \leq 2$. ■

[6] This is one of these "trivial" claims that might give the reader a pause, so here is a formal proof. Pick a random point p uniformly inside the ball $\mathbf{b}^n$. Let ψ be the probability that $\mathsf{p} \in \widehat{A}$. Clearly, $\mathrm{vol}\left(\widehat{A}\right) = \psi\mathrm{vol}(\mathbf{b}^n)$. So, consider the normalized point $\mathsf{q} = \mathsf{p}/\|\mathsf{p}\|$. Clearly, $\mathsf{p} \in \widehat{A}$ if and only if $\mathsf{q} \in A$, by the definition of $\widehat{A}$. Thus, $\mu\left(\widehat{A}\right) = \mathrm{vol}\left(\widehat{A}\right)/\mathrm{vol}(\mathbf{b}^n) = \psi = \mathbf{Pr}\left[\mathsf{p} \in \widehat{A}\right] = \mathbf{Pr}[\mathsf{q} \in A] = \mathbf{Pr}[A]$, since q has a uniform distribution on the hypersphere by assumption.

19.3. Concentration of Lipschitz Functions

Consider a function $f : \mathbb{S}^{(n-1)} \to \mathbb{R}$, and imagine that we have a probability density function defined over the sphere. Let $\mathbf{Pr}[f \leq t] = \mathbf{Pr}\Big[\Big\{x \in S^{n-1} \;\Big|\; f(x) \leq t\Big\}\Big]$. We define the *median* of f, denoted by $\mathrm{med}(f)$, to be the $\sup t$, such that $\mathbf{Pr}[f \leq t] \leq 1/2$.

We define $\mathbf{Pr}[f < \mathrm{med}(f)] = \sup_{x<\mathrm{med}(f)} \mathbf{Pr}[f \leq x]$. The following is obvious but (in fact) requires a formal proof.

Lemma 19.3.1. *We have* $\mathbf{Pr}[f < \mathrm{med}(f)] \leq 1/2$ *and* $\mathbf{Pr}[f > \mathrm{med}(f)] \leq 1/2$.

Proof: Since $\bigcup_{k\geq 1}(-\infty, \mathrm{med}(f) - 1/k] = (-\infty, \mathrm{med}(f))$, we have

$$\mathbf{Pr}[f < \mathrm{med}(f)] \quad = \quad \sup_{k\geq 1} \mathbf{Pr}\left[f \leq \mathrm{med}(f) - \frac{1}{k}\right] \leq \sup_{k\geq 1} \frac{1}{2} = \frac{1}{2}.$$

The second claim follows by a symmetric argument. ■

Definition 19.3.2 (*c*-***Lipschitz***). A function $f : A \to B$ is *c*-***Lipschitz*** if, for any $x, y \in A$, we have $\|f(x) - f(y)\| \leq c\,\|x - y\|$.

Theorem 19.3.3 (Lévy's Lemma). *Let $f : \mathbb{S}^{(n-1)} \to \mathbb{R}$ be 1-Lipschitz. Then for all $t \in [0, 1]$,*

$$\mathbf{Pr}[f > \mathrm{med}(f) + t] \leq 2\exp\left(-t^2 n/2\right) \quad \textit{and} \quad \mathbf{Pr}[f < \mathrm{med}(f) - t] \leq 2\exp\left(-t^2 n/2\right).$$

Proof: We prove only the first inequality, the second follows by symmetry. Let

$$A = \left\{x \in \mathbb{S}^{(n-1)} \;\middle|\; f(x) \leq \mathrm{med}(f)\right\}.$$

By Lemma 19.3.1, we have $\mathbf{Pr}[A] \geq 1/2$. Consider a point $x \in A_t$, where A_t is as defined in Theorem 19.2.1. Let $\mathsf{nn}(x)$ be the nearest point in A to x. We have by definition that $\|x - \mathsf{nn}(x)\| \leq t$. As such, since f is 1-Lipschitz and $\mathsf{nn}(x) \in A$, we have that

$$f(x) \leq f(\mathsf{nn}(x)) + \|\mathsf{nn}(x) - x\| \leq \mathrm{med}(f) + t.$$

Thus, by Theorem 19.2.1, we get $\mathbf{Pr}[f > \mathrm{med}(f) + t] \leq 1 - \mathbf{Pr}[A_t] \leq 2\exp\left(-t^2 n/2\right)$. ■

19.4. The Johnson-Lindenstrauss Lemma

Lemma 19.4.1. *For a unit vector $x \in \mathbb{S}^{(n-1)}$, let*

$$f(x) = \sqrt{x_1^2 + x_2^2 + \cdots + x_k^2}$$

be the length of the projection of x into the subspace formed by the first k coordinates. Let x be a vector randomly chosen with uniform distribution from $\mathbb{S}^{(n-1)}$. Then $f(x)$ is sharply concentrated. Namely, there exists $m = m(n, k)$ such that

$$\mathbf{Pr}[f(x) \geq m + t] \leq 2\exp(-t^2 n/2) \quad \textit{and} \quad \mathbf{Pr}[f(x) \leq m - t] \leq 2\exp(-t^2 n/2),$$

for any $t \in [0, 1]$. Furthermore, for $k \geq 10 \ln n$, we have $m \geq \frac{1}{2}\sqrt{k/n}$.

Proof: The orthogonal projection $p : \mathbb{R}^n \to \mathbb{R}^k$ given by $p(x_1, \ldots, x_n) = (x_1, \ldots, x_k)$ is 1-Lipschitz (since projections can only shrink distances, see Exercise 19.6.4). As such, $f(x) = \|p(x)\|$ is 1-Lipschitz, since for any x, y we have

$$|f(x) - f(y)| = |\|p(x)\| - \|p(y)\|| \leq \|p(x) - p(y)\| \leq \|x - y\|,$$

by the triangle inequality and since p is 1-Lipschitz. Theorem 19.3.3 (i.e., Lévy's lemma) gives the required tail estimate with $m = \operatorname{med}(f)$.

Thus, we only need to prove the lower bound on m. For a random $x = (x_1, \ldots, x_n) \in \mathbb{S}^{(n-1)}$, we have $\mathbf{E}\left[\|x\|^2\right] = 1$. By linearity of expectations, and symmetry, we have $1 = \mathbf{E}\left[\|x\|^2\right] = \mathbf{E}\left[\sum_{i=1}^n x_i^2\right] = \sum_{i=1}^n \mathbf{E}\left[x_i^2\right] = n\,\mathbf{E}\left[x_j^2\right]$, for any $1 \leq j \leq n$. Thus, $\mathbf{E}\left[x_j^2\right] = 1/n$, for $j = 1, \ldots, n$. Thus,

$$\mathbf{E}\left[(f(x))^2\right] = \mathbf{E}\left[\sum_{i=1}^k x_i^2\right] = \sum_{i=1}^k \mathbf{E}[x_i] = \frac{k}{n},$$

by linearity of expectation.

We next use that f is concentrated, to show that f^2 is also relatively concentrated. For any $t \geq 0$, we have

$$\frac{k}{n} = \mathbf{E}\left[f^2\right] \leq \mathbf{Pr}[f \leq m + t]\,(m + t)^2 + \mathbf{Pr}[f \geq m + t] \cdot 1 \leq 1 \cdot (m + t)^2 + 2\exp(-t^2 n/2),$$

since $f(x) \leq 1$, for any $x \in \mathbb{S}^{(n-1)}$. Let $t = \sqrt{k/5n}$. Since $k \geq 10 \ln n$, we have that $2\exp(-t^2 n/2) \leq 2/n$. We get that

$$\frac{k}{n} \leq \left(m + \sqrt{k/5n}\right)^2 + 2/n.$$

Implying that $\sqrt{(k-2)/n} \leq m + \sqrt{k/5n}$, which in turn implies that $m \geq \sqrt{(k-2)/n} - \sqrt{k/5n} \geq \frac{1}{2}\sqrt{k/n}$. ■

Next, we would like to argue that given a fixed vector, projecting it down into a random k-dimensional subspace results in a random vector such that its length is highly concentrated. This would imply that we can do dimension reduction and still preserve distances between points that we care about.

To this end, we would like to flip Lemma 19.4.1 around. Instead of randomly picking a point and projecting it down to the first k-dimensional space, we would like x to be fixed, and randomly pick the k-dimensional subspace we project into. However, we need to pick this random k-dimensional space carefully. Indeed, if we rotate this random subspace, by a transformation T, so that it occupies the first k dimensions, then the point $T(x)$ needs to be uniformly distributed on the hypersphere if we want to use Lemma 19.4.1.

As such, we would like to randomly pick a rotation of $\mathbb{R}^n$. This maps the standard orthonormal basis into a randomly rotated orthonormal space. Taking the subspace spanned by the first k vectors of the rotated basis results in a k-dimensional random subspace. Such a rotation is an orthonormal matrix with determinant 1. We can generate such a matrix, by randomly picking a vector $e_1 \in \mathbb{S}^{(n-1)}$. Next, we set e_1 as the first column of our rotation matrix, and generate the other $n - 1$ columns, by generating recursively $n - 1$ orthonormal vectors in the space orthogonal to e_1.

Remark 19.4.2 (Generating a random point on the sphere.). At this point, the reader might wonder how do we pick a point uniformly from the unit hypersphere. The idea is to pick a point from the multi-dimensional normal distribution $N^n(0, 1)$, and normalizing it to have length 1. Since the multi-dimensional normal distribution has the density function

$$(2\pi)^{-n/2} \exp\left(-(x_1^2 + x_2^2 + \cdots + x_n^2)/2\right),$$

which is symmetric (i.e., all the points in distance r from the origin have the same distribution), it follows that this indeed generates a point randomly and uniformly on $\mathbb{S}^{(n-1)}$.

Generating a vector with multi-dimensional normal distribution, is no more than picking each coordinate according to the normal distribution, see Lemma 19.7.1$_{\text{p13}}$. Given a source of random numbers according to the uniform distribution, this can be done using a $O(1)$ computations per coordinate, using the Box-Muller transformation [BM58]. Overall, each random vector can be generated in $O(n)$ time.

Since projecting down n-dimensional normal distribution to the lower dimensional space yields a normal distribution, it follows that generating a random projection, is no more than randomly picking n vectors according to the multidimensional normal distribution $v_1, \ldots, v_n$. Then, we orthonormalize them, using Graham-Schmidt, where $\widehat{v_1} = v_1 / \|v_1\|$, and $\widehat{v_i}$ is the normalized vector of $v_i - w_i$, where w_i is the projection of v_i to the space spanned by $v_1, \ldots, v_{i-1}$.

Taking those vectors as columns of a matrix, generates a matrix A, with determinant either 1 or -1. We multiply one of the vectors by -1 if the determinant is -1. The resulting matrix is a random rotation matrix.

We can now restate Lemma 19.4.1 in the setting where the vector is fixed and the projection is into a random subspace.

Lemma 19.4.3. *Let $x \in \mathbb{S}^{(n-1)}$ be an arbitrary unit vector, and consider a random k dimensional subspace $\mathcal{F}$, and let $f(x)$ be the length of the projection of x into $\mathcal{F}$. Then, there exists $m = m(n,k)$ such that*

$$\mathbf{Pr}[f(x) \geq m + t] \leq 2\exp(-t^2 n/2) \quad \text{and} \quad \mathbf{Pr}[f(x) \leq m - t] \leq 2\exp(-t^2 n/2),$$

for any $t \in [0,1]$. Furthermore, for $k \geq 10 \ln n$, we have $m \geq \frac{1}{2}\sqrt{k/n}$.

Proof: Let v_i be the ith orthonormal vector having 1 at the ith coordinate. Let M be a random translation of space generated as described above. Clearly, for arbitrary fixed unit vector x, the vector $\mathsf{M}x$ is distributed uniformly on the sphere. Now, the ith column of the matrix M is the random vector e_i, and $\mathsf{M}^T v_i = e_i$. As such, we have

$$\langle \mathsf{M}x, v_i \rangle = (\mathsf{M}x)^T v_i = x^T \mathsf{M}^T v_i = x^T e_i = \langle x, e_i \rangle .$$

In particular, treating $\mathsf{M}x$ as a random vector, and projecting it on the first k coordinates, we have that

$$f(x) = \sqrt{\sum_{i=1}^{k} \langle \mathsf{M}x, v_i \rangle^2} = \sqrt{\sum_{i=1}^{k} \langle x, e_i \rangle^2}.$$

But $e_1, \ldots, e_k$ is just an orthonormal basis of a random k-dimensional subspace. As such, the expression on the right is the length of the projection of x into a k-dimensional random subspace. As such, the length of the projection of x into a random k-dimensional subspace has exactly the same distribution as the length of the projection of a random vector into the first k coordinates. The claim now follows by Lemma 19.4.1. ■

Definition 19.4.4. The mapping $f : \mathbb{R}^n \to \mathbb{R}^k$ is called ***K-bi-Lipschitz*** for a subset $X \subseteq \mathbb{R}^n$ if there exists a constant $c > 0$ such that

$$cK^{-1} \cdot \|\mathsf{p} - \mathsf{q}\| \leq \|f(\mathsf{p}) - f(\mathsf{q})\| \leq c \cdot \|\mathsf{p} - \mathsf{q}\| ,$$

for all $\mathsf{p}, \mathsf{q} \in X$.

The least K for which f is K-bi-Lipschitz is called the *distortion* of f, and is denoted dist(f). We will refer to f as a ***K-embedding*** of X.

Remark 19.4.5. Let $X \subseteq \mathbb{R}^m$ be a set of n points, where m potentially might be much larger than n. Observe, that in this case, since we only care about the inter-point distances of points in X, we can consider X to be a set of points lying in the affine subspace $\mathcal{F}$ spanned by the points of X. Note, that this subspace has dimension $n-1$. As such, each point of X be interpreted as $n-1$ dimensional point in $\mathcal{F}$. Namely, we can assume, for our purposes, that the set of n points in Euclidean space we care about lies in $\mathbb{R}^n$ (in fact, $\mathbb{R}^{n-1}$).

Note, that if $m < n$ we can always pad all the coordinates of the points of X by zeros, such that the resulting point set lies in $\mathbb{R}^n$.

Theorem 19.4.6 (Johnson-Lindenstrauss lemma.). *Let X be an n-point set in a Euclidean space, and let $\varepsilon \in (0,1]$ be given. Then there exists a $(1+\varepsilon)$-embedding of X into $\mathbb{R}^k$, where $k = O(\varepsilon^{-2}\log n)$.*

Proof: By Remark 19.4.5, we can assume that $X \subseteq \mathbb{R}^n$. Let $k = 200\varepsilon^{-2}\ln n$. Assume $k < n$, and let $\mathcal{F}$ be a random k-dimensional linear subspace of $\mathbb{R}^n$. Let $P_{\mathcal{F}} : \mathbb{R}^n \to \mathcal{F}$ be the orthogonal projection operator of $\mathbb{R}^n$ into $\mathcal{F}$. Let m be the number around which $\|P_{\mathcal{F}}(x)\|$ is concentrated, for $x \in \mathbb{S}^{(n-1)}$, as in Lemma 19.4.3.

Fix two points $x, y \in \mathbb{R}^n$, we prove that

$$\left(1-\frac{\varepsilon}{3}\right)m\|x-y\| \le \|P_{\mathcal{F}}(x)-P_{\mathcal{F}}(y)\| \le \left(1+\frac{\varepsilon}{3}\right)m\|x-y\|$$

holds with probability $\ge 1-n^{-2}$. Since there are $\binom{n}{2}$ pairs of points in X, it follows that with constant probability (say $> 1/3$) this holds for all pairs of points of X. In such a case, the mapping p is D-embedding of X into $\mathbb{R}^k$ with $D \le \frac{1+\varepsilon/3}{1-\varepsilon/3} \le 1+\varepsilon$, for $\varepsilon \le 1$.

Let $u = x - y$, we have $P_{\mathcal{F}}(u) = P_{\mathcal{F}}(x) - P_{\mathcal{F}}(y)$ since $P_{\mathcal{F}}(\cdot)$ is a linear operator. Thus, the condition becomes $\left(1-\frac{\varepsilon}{3}\right)m\|u\| \le \|P_{\mathcal{F}}(u)\| \le \left(1+\frac{\varepsilon}{3}\right)m\|u\|$. Again, since projection is a linear operator, for any $\alpha > 0$, the condition is equivalent to

$$\left(1-\tfrac{\varepsilon}{3}\right)m\|\alpha u\| \le \|P_{\mathcal{F}}(\alpha u)\| \le \left(1+\tfrac{\varepsilon}{3}\right)m\|\alpha u\|.$$

As such, we can assume that $\|u\| = 1$ by picking $\alpha = 1/\|u\|$. Namely, we need to show that

$$\left|\|P_{\mathcal{F}}(u)\| - m\right| \le \frac{\varepsilon}{3}m.$$

Let $f(u) = \|P_{\mathcal{F}}(u)\|$. By Lemma 19.4.1 (exchanging the random space with the random vector), for $t = \varepsilon m/3$, we have that the probability that this does not hold is bounded by

$$\mathbf{Pr}[|f(u)-m| \ge t] \le 4\exp\left(-\frac{t^2 n}{2}\right) = 4\exp\left(\frac{-\varepsilon^2 m^2 n}{18}\right) \le 4\exp\left(-\frac{\varepsilon^2 k}{72}\right) < n^{-2},$$

since $m \ge \frac{1}{2}\sqrt{k/n}$ and $k = 200\varepsilon^{-2}\ln n$. ∎

19.5. Bibliographical notes

Our presentation follows Matoušek [Mat02]. The Brunn-Minkowski inequality is a powerful inequality which is widely used in mathematics. A nice survey of this inequality and its applications is provided by Gardner [Gar02]. Gardner says: "In a sea of mathematics, the Brunn-Minkowski inequality appears like an octopus, tentacles reaching far and wide, its shape and color changing as it roams from one area to the next." However, Gardner is careful in claiming that the Brunn-Minkowski inequality is one of the most powerful inequalities

in mathematics since as a wit put it "the most powerful inequality is $x^2 \geq 0$, since all inequalities are in some sense equivalent to it."

A striking application of the Brunn-Minkowski inequality is the proof that in any partial ordering of n elements, there is a single comparison that knowing its result, reduces the number of linear extensions that are consistent with the partial ordering, by a constant fraction. This immediately implies (the uninteresting result) that one can sort n elements in $O(n \log n)$ comparisons. More interestingly, it implies that if there are m linear extensions of the current partial ordering, we can *always* sort it using $O(\log m)$ comparisons. A nice exposition of this surprising result is provided by Matoušek [Mat02, Section 12.3].

There are several alternative proofs of the JL lemma, see [IM98] and [DG03]. Interestingly, it is enough to pick each entry in the dimension reducing matrix randomly out of $-1, 0, 1$. This requires a more involved proof [Ach01]. This is useful when one cares about storing this dimension reduction transformation efficiently.

Magen [Mag07] observed that the JL lemma preserves angles, and in fact can be used to preserve any "k dimensional angle", by projecting down to dimension $O(k\varepsilon^{-2} \log n)$. In particular, Exercise 19.6.5 is taken from there.

In fact, the random embedding preserves much more structure than just distances between points. It preserves the structure and distances of surfaces as long as they are low dimensional and "well behaved", see [AHY07] for some results in this direction.

Dimension reduction is crucial in learning, AI, databases, etc. One common technique that is being used in practice is to do PCA (i.e., principal component analysis) and take the first few main axes. Other techniques include independent component analysis, and MDS (multidimensional scaling). MDS tries to embed points from high dimensions into low dimension ($d = 2$ or 3), while preserving some properties. Theoretically, dimension reduction into really low dimensions is hopeless, as the distortion in the worst case is $\Omega(n^{1/(k-1)})$, if k is the target dimension [Mat90].

19.6. Exercises

Exercise 19.6.1 (Boxes can be separated.). (Easy.) Let A and B be two axis-parallel boxes that are interior disjoint. Prove that there is always an axis-parallel hyperplane that separates the interior of the two boxes.

Exercise 19.6.2 (Brunn-Minkowski inequality slight extension.). Prove the following.

> **Corollary 19.6.3.** *For A and B compact sets in $\mathbb{R}^n$, we have for any $\lambda \in [0, 1]$ that* $\mathrm{vol}(\lambda A + (1-\lambda)B) \geq \mathrm{vol}(A)^{\lambda}\mathrm{vol}(B)^{1-\lambda}$.

Exercise 19.6.4 (Projections are contractions.). (Easy.) Let $\mathcal{F}$ be a k-dimensional affine subspace, and let $P_{\mathcal{F}} : \mathbb{R}^d \to \mathcal{F}$ be the projection that maps every point $x \in \mathbb{R}^d$ to its nearest neighbor on $\mathcal{F}$. Prove that p is a contraction (i.e., 1-Lipschitz). Namely, for any $\mathsf{p}, \mathsf{q} \in \mathbb{R}^d$, it holds that $\|P_{\mathcal{F}}(\mathsf{p}) - P_{\mathcal{F}}(\mathsf{q})\| \leq \|\mathsf{p} - \mathsf{q}\|$.

Exercise 19.6.5 (JL Lemma works for angles.). Show that the Johnson-Lindenstrauss lemma also $(1 \pm \varepsilon)$-preserves angles among triples of points of P (you might need to increase the target dimension however by a constant factor). *[For every angle, construct a equilateral triangle that its edges are being preserved by the projection (add the vertices of those triangles [conceptually] to the point set being embedded). Argue, that this implies that the angle is being preserved.]*

19.7. Miscellaneous

Lemma 19.7.1. *(A) The multidimensional normal distribution is symmetric; that is, for any two points* $\mathsf{p}, \mathsf{q} \in \mathbb{R}^d$ *such that* $\|\mathsf{p}\| = \|\mathsf{q}\|$ *we have that* $g(\mathsf{p}) = g(\mathsf{q})$*, where* $g(\cdot)$ *is the density function of the multidimensional normal distribution* $\mathbf{N}^d$*.*

(B) The projection of the normal distribution on any direction is a one dimensional normal distribution.

(C) Picking d variables $X_1, \ldots, X_d$ *using one dimensional normal distribution* $\mathbf{N}$ *results in a point* $(X_1, \ldots, X_d)$ *that has multidimensional normal distribution* $\mathbf{N}^d$*.*

Bibliography

[Ach01] D. Achlioptas. Database-friendly random projections. In *Proc. 20th ACM Sympos. Principles Database Syst.* (PODS), pages 274–281, 2001.

[AHY07] P. Agarwal, S. Har-Peled, and H. Yu. Embeddings of surfaces, curves, and moving points in Euclidean space. In *Proc. 23rd Annu. Sympos. Comput. Geom.* (SoCG), pages 381–389, 2007.

[BM58] G. E.P. Box and M. E. Muller. A note on the generation of random normal deviates. *Ann. Math. Stat.*, 28:610–611, 1958.

[Car76] L. Carroll. The hunting of the snark, 1876.

[DG03] S. Dasgupta and A. Gupta. An elementary proof of a theorem of Johnson and Lindenstrauss. *Rand. Struct. Alg.*, 22(3):60–65, 2003.

[Gar02] R. J. Gardner. The Brunn-Minkowski inequality. *Bull. Amer. Math. Soc.*, 39:355–405, 2002.

[IM98] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. 30th Annu. ACM Sympos. Theory Comput.* (STOC), pages 604–613, 1998.

[Mag07] A. Magen. Dimensionality reductions in ℓ_2 that preserve volumes and distance to affine spaces. *Discrete Comput. Geom.*, 38(1):139–153, 2007.

[Mat90] J. Matoušek. Bi-Lipschitz embeddings into low-dimensional Euclidean spaces. *Comment. Math. Univ. Carolinae*, 31:589–600, 1990.

[Mat02] J. Matoušek. *Lectures on Discrete Geometry*, volume 212 of *Grad. Text in Math.* Springer, 2002.