

11

Dimension Reduction and the JL Lemma

For a set of n points $\{x_1, x_2, \dots, x_n\}$ in \mathbb{R}^D , can we map them into some lower dimensional space \mathbb{R}^k and still maintain the Euclidean distances between them? We can always take $k \leq n - 1$, since any set of n points lies on a $n - 1$ -dimensional subspace. And this is (existentially) tight, e.g., if $x_2 - x_1, x_3 - x_1, \dots, x_n - x_1$ are all orthogonal vectors.

But what if we were fine with distances being approximately preserved? There can only be k orthogonal unit vectors in \mathbb{R}^k , but there are as many as $\exp(c\epsilon^2 k)$ unit vectors which are ϵ -orthogonal—i.e., whose mutual inner products all lie in $[-\epsilon, \epsilon]$. Near-orthogonality allows us to pack exponentially more vectors! (Indeed, we will see this in a homework exercise.)

This near-orthogonality of the unit vectors means that distances are also approximately preserved. Indeed, for any two $a, b \in \mathbb{R}^k$,

$$\|a - b\|_2^2 = \langle a - b, a - b \rangle = \langle a, a \rangle + \langle b, b \rangle - 2\langle a, b \rangle = \|a\|_2^2 + \|b\|_2^2 - 2\langle a, b \rangle,$$

so the squared Euclidean distance between any pair of the points defined by these ϵ -orthogonal vectors falls in the range $2(1 \pm \epsilon)$. So, if we wanted n points at exactly the same (Euclidean) distance from each other, we would need $n - 1$ dimensions. (Think of a triangle in 2-dims.) But if we wanted to pack in n points which were at distance $(1 \pm \epsilon)$ from each other, we could pack them into

$$k = O\left(\frac{\log n}{\epsilon^2}\right)$$

dimensions.

Having $n \geq \exp(c\epsilon^2 k)$ vectors in d dimensions means the dimension is $k = O(\log n / \epsilon^2)$.

11.1 The Johnson Lindenstrauss lemma

The Johnson Lindenstrauss “flattening” lemma says that such a claim is true not just for equidistant points, but for any set of n points in Euclidean space:

Lemma 11.1. Let $\varepsilon \in (0, 1/2)$. Given any set of points $X = \{x_1, x_2, \dots, x_n\}$ in \mathbb{R}^D , there exists a map $A : \mathbb{R}^D \rightarrow \mathbb{R}^k$ with $k = O\left(\frac{\log n}{\varepsilon^2}\right)$ such that

$$1 - \varepsilon \leq \frac{\|A(x_i) - A(x_j)\|_2^2}{\|x_i - x_j\|_2^2} \leq 1 + \varepsilon.$$

Moreover, such a map can be computed in expected $\text{poly}(n, D, 1/\varepsilon)$ time.

Note that the target dimension k is independent of the original dimension D , and depends only on the number of points n and the accuracy parameter ε .

It is not difficult to show that we need at least $\Omega(\log n)$ dimensions in such a result, using a packing argument. Noga Alon showed a lower bound of $\Omega\left(\frac{\log n}{\varepsilon^2 \log 1/\varepsilon}\right)$, and then Kasper Green Larson and Jelani Nelson showed a tight and matching lower bound of $\Omega\left(\frac{\log n}{\varepsilon^2}\right)$ dimensions for any dimensionality reduction scheme from n dimensions that preserves pairwise distances.

The JL Lemma was first considered in the area of metric embeddings, for applications like fast near-neighbor searching; today we use it to speed up algorithms for problems like spectral sparsification of graphs, and solving linear programs fast.

11.2 The Construction

The JL lemma is pretty surprising, but the construction of the map is perhaps even more surprising: it is a super-simple randomized construction. Let M be a $k \times D$ matrix, such that every entry of M is filled with an i.i.d. draw from a standard normal $N(0, 1)$ distribution (a.k.a. the “Gaussian” distribution). For $x \in \mathbb{R}^D$, define

$$A(x) = \frac{1}{\sqrt{k}} Mx.$$

That’s it. You hit the vector x with a Gaussian matrix M , and scale it down by \sqrt{k} . That’s the map A .

Since $A(x)$ is a linear map and satisfies $\alpha A(x) + \beta A(y) = A(\alpha x + \beta y)$, it is enough to show the following lemma:

Lemma 11.2. [Distributional Johnson-Lindenstrauss] Let $\varepsilon \in (0, 1/2)$. If A is constructed as above with $k = c\varepsilon^{-2} \log \delta^{-1}$, and $x \in \mathbb{R}^D$ is a unit vector, then

$$\Pr[\|A(x)\|_2^2 \in 1 \pm \varepsilon] \geq 1 - \delta.$$

To prove Lemma 11.1, set $\delta = 1/n^2$, and hence $k = O(\varepsilon^{-2} \log n)$.

Now for each $x_i, x_j \in X$, use linearity of $A(\cdot)$ to infer

$$\frac{\|A(x_i) - A(x_j)\|^2}{\|x_i - x_j\|^2} = \frac{\|A(x_i - x_j)\|^2}{\|x_i - x_j\|^2} = \|A(v_{ij})\|^2 \in (1 \pm \varepsilon)$$

Given n points with Euclidean distances in $(1 \pm \varepsilon)$, the balls of radius $\frac{1-\varepsilon}{2}$ around these points must be mutually disjoint, by the minimum distance, and they are contained within a ball of radius $(1 + \varepsilon) + \frac{1-\varepsilon}{2}$ around x_0 . Since volumes of balls in \mathbb{R}^k of radius r behave like $c_k r^k$, we have

$$n \cdot c_k \left(\frac{1-\varepsilon}{2}\right)^k \leq c_k \left(\frac{3+\varepsilon}{2}\right)^k$$

or $k \geq \Omega(\log n)$ for $\varepsilon \leq 1/2$.

Alon (2003)

Larson and Nelson (2017)

with probability at least $1 - 1/n^2$, where v_{ij} is the unit vector in the direction of $x_i - x_j$. By a union bound, all $\binom{n}{2}$ pairs of distances in $\binom{X}{2}$ are maintained with probability at least $1 - \binom{n}{2} \frac{1}{n^2} \geq 1/2$. A few comments about this construction:

- The above proof shows not only the existence of a good map, we also get that a random map as above works with constant probability! In other words, a Monte-Carlo randomized algorithm for dimension reduction. (Since we can efficiently check that the distances are preserved to within the prescribed bounds, we can convert this into a Las Vegas algorithm.) Or we can also get deterministic algorithms: see [here](#).
- The algorithm (at least the Monte Carlo version) is *data-oblivious*: it does not even look at the set of points X : it works for any set X with high probability. Hence, we can pick this map A before the points in X arrive.

11.3 Intuition for the Distributional JL Lemma

Let us recall some basic facts about Gaussian distributions. The probability density function for the Gaussian $N(\mu, \sigma^2)$ is

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}.$$

We also use the following; the proof just needs some elbow grease.

Proposition 11.3. *If $G_1 \sim N(\mu_1, \sigma_1^2)$ and $G_2 \sim N(\mu_2, \sigma_2^2)$ are independent, then for $c \in \mathbb{R}$,*

$$c G_1 \sim N(c\mu_1, c^2 \sigma_1^2) \tag{11.1}$$

$$G_1 + G_2 \sim N(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2). \tag{11.2}$$

Now, here's the main idea in the proof of Lemma 11.2. Imagine that the vector x is the elementary unit vector $e_1 = (1, 0, \dots, 0)$. Then $M e_1$ is just the first column of M , which is a vector with independent and identical Gaussian values.

$$M e_1 = \begin{bmatrix} G_{1,1} & G_{1,2} & \cdots & G_{1,D} \\ G_{2,1} & G_{2,2} & \cdots & G_{2,D} \\ \vdots & \vdots & \ddots & \vdots \\ G_{k,1} & G_{k,2} & \cdots & G_{k,D} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} G_{1,1} \\ G_{2,1} \\ \vdots \\ G_{k,1} \end{bmatrix}.$$

$A(x)$ is a scaling-down of this vector by \sqrt{k} : every entry in this random vector $A(x) = A(e_1)$ is distributed as

$$1/\sqrt{k} \cdot N(0, 1) = N(0, 1/k) \quad (\text{by (11.1)}).$$

The fact that the means and the variances take on the claimed values should not be surprising; this is true for all r.v.s. The surprising part is that the resulting variables are also Gaussians.

Thus, the expected squared length of $A(x) = A(e_1)$ is

$$\mathbb{E} \left[\|A(x)\|^2 \right] = \mathbb{E} \left[\sum_{i=1}^k A(x)_i^2 \right] = \sum_{i=1}^k \mathbb{E} \left[A(x)_i^2 \right] = \sum_{i=1}^k \frac{1}{k} = 1.$$

So the expectation of $\|A(x)\|^2$ is 1; the heart is in the right place!

Now to show that $\|A(x)\|^2$ does not deviate too much from the mean—i.e., to show a concentration result. Indeed, $\|A(x)\|^2$ is a sum of independent $N(0, 1/k)^2$ random variables, so if these $N(0, 1/k)^2$ variables were bounded, we would be done by the Chernoff bounds of the previous chapter. Sadly, they are not. However, their tails are fairly “thin”, so if we squint hard enough, these random variables can be viewed as “pretty much bounded”, and the Chernoff bounds can be used.

Of course this is very vague and imprecise. Indeed, the Laplace distribution with density function $f(x) \propto e^{-\lambda|x|}$ for $x \in \mathbb{R}$ also has pretty thin tails—“exponential tails”. But using a matrix with Laplace entries does not work the same, no matter how hard we squint. It turns out you need the entries of M , the matrix used to define $A(x)$, to have “sub-Gaussian tails”. The Gaussian entries have precisely this property.

We now make all this precise, and also remove the assumption that the vector $x = e_1$. In fact, we do this in two ways.

1. First we give a proof via a direct calculation: it has several steps, but each step is elementary, and you are mostly following your nose.
2. The second proof uses the notion of sub-Gaussian random variables from , and builds some general machinery for concentration bounds.

11.4 A Direct Proof of Lemma 11.2

Recall that we want to argue about the squared length of $A(x) \in \mathbb{R}^k$, where $A(x) = \frac{1}{\sqrt{k}} Mx$, and x is a unit vector. First, let’s understand what the expected length of $A(x)$ is, and then we will show concentration about the mean.

Lemma 11.4. *Suppose the entries of M are independent random variables, with mean zero and unit variance. Then for unit vector $x \in \mathbb{R}^D$,*

$$\mathbb{E}[\|A(x)\|^2] = \|x\|^2.$$

Proof. Each entry of the vector Mx is the inner product of x with a vector with independent zero mean and unit variance random

If G has mean μ and variance σ^2 , then $\mathbb{E}[G^2] = \text{Var}[G] + \mathbb{E}[G]^2 = \sigma^2 + \mu^2$.

variables, and so is itself a random variable with zero mean and variance $\sum_i x_i^2 = 1$. This means that for any entry $i \in [k]$,

$$\mathbb{E}[(Mx)_i^2] = \text{Var}(Mx) + \mathbb{E}[(Mx)_i]^2 = 1.$$

Now $\mathbb{E}[\|A(x)\|^2] = \frac{1}{k} \sum_{i=1}^k \mathbb{E}[(Mx)_i^2] = 1 = \|x\|^2$. \square

Observe that did not use the fact that the matrix entries were Gaussians. We will use it for the concentration bound, which we show next.

11.4.1 Concentration about the Mean

Using that each entry of M is an independent $N(0, 1)$ r.v., we can use Proposition 11.3 to infer that $(Mx)_i \sim N(0, x_1^2 + x_2^2 + \dots + x_D^2) = N(0, 1)$. So, each of the k coordinates of Mx behaves just like an independent Gaussian! For brevity, define

$$Z := \|A(z)\|^2 = \sum_{i=1}^k \frac{1}{k} (Mx)_i^2,$$

so Z is the *average* of the *squares* of a collection of k independent $N(0, 1)$ r.v.s.

Next we show that Z does not deviate too much from 1. Since Z is the sum of a bunch of independent and identical random variables, let's start down the usual path for a Chernoff bound, for the upper tail, say:

$$\Pr[Z \geq 1 + \varepsilon] \leq \Pr[e^{tZ} \geq e^{t(1+\varepsilon)}] \leq \mathbb{E}[e^{tZ}] / e^{t(1+\varepsilon)} \quad (11.3)$$

$$= \prod_i \left(\mathbb{E}[e^{tG^2}] / e^{t(1+\varepsilon)} \right) \quad (11.4)$$

for every $t > 0$, where $G \sim N(0, 1)$. Now $\mathbb{E}[e^{tG^2}]$, the moment-generating function for G^2 is easy to calculate for $t < 1/2$:

$$\frac{1}{\sqrt{2\pi}} \int_{g \in \mathbb{R}} e^{tg^2} e^{-g^2/2} dg = \frac{1}{\sqrt{2\pi}} \int_{z \in \mathbb{R}} e^{-z^2/2} \frac{dz}{\sqrt{1-2t}} = \frac{1}{\sqrt{1-2t}}. \quad (11.5)$$

Plugging back into (11.4), the bound on the upper tail shows that for all $t \in (0, 1/2)$,

$$\Pr[Z \geq (1 + \varepsilon)] \leq \left(\frac{1}{e^{t(1+\varepsilon)} \sqrt{1-2t}} \right)^k.$$

The easy way out is to observe that the squares of Gaussians are [chi-squared](#) r.v.s, the sum of k of them is χ^2 with k degrees of freedom, and the internet conveniently has [tail bounds](#) for these things. But even if you don't recall these facts, and don't have internet connectivity and cannot check Wikipedia, it is not that difficult to prove from scratch.

Let's just focus on part of this expression:

$$\left(\frac{1}{e^t \sqrt{1-2t}} \right) = \exp \left(-t - \frac{1}{2} \log(1-2t) \right) \quad (11.6)$$

$$= \exp \left((2t)^2/4 + (2t)^3/6 + \dots \right) \quad (11.7)$$

$$\leq \exp \left(t^2(1+2t+2t^2+\dots) \right) \quad (11.8)$$

$$= \exp(t^2/(1-2t)).$$

Plugging this back, we get

$$\begin{aligned} \Pr[Z \geq (1+\varepsilon)] &\leq \left(\frac{1}{e^{t(1+\varepsilon)} \sqrt{1-2t}} \right)^k \\ &\leq \exp(kt^2/(1-2t) - kte) \leq e^{-k\varepsilon^2/8}, \end{aligned}$$

if we set $t = \varepsilon/4$ and use the fact that $1-2t \geq 1/2$ for $\varepsilon \leq 1/2$. (Note: this setting of t also satisfies $t \in (0, 1/2)$, which we needed from our previous calculations.)

Almost done: let's take stock of the situation. We observed that $\|A(x)\|_2^2$ was distributed like an average of squares of Gaussians, and by a Chernoff-like calculation we proved that

$$\Pr[\|A(x)\|_2^2 > 1+\varepsilon] \leq \exp(-k\varepsilon^2/8) \leq \delta/2$$

for $k = \frac{8}{\varepsilon^2} \ln \frac{2}{\delta}$. A similar calculation bounds the lower tail, and finishes the proof of Lemma 11.2.

The JL Lemma was first proved by Bill Johnson and Joram Lindenstrauss. There have been several proofs after theirs, usually trying to tighten their results, or simplify the algorithm/proof (see citations in some of the newer papers): the proof above is some combinations of those by [Piotr Indyk and Rajeev Motwani](#), and [Sanjoy Dasgupta and myself](#).

[Johnson and Lindenstrauss \(1982\)](#)

[Indyk and Motwani \(1998\)](#)

[Dasgupta and Gupta \(2004\)](#)

11.5 Introducing Subgaussian Random Variables

It turns out that the proof of Lemma 11.2 is a bit cleaner (with fewer calculations) if we use the abstraction provided by the generic Chernoff bound from last lecture, and the notion of subGaussian random variables which we introduce next. This abstraction will also allow us to extend the result to JL matrices having i.i.d. entries from other distributions, e.g., where each $M_{ij} \in_R \{-1, +1\}$.

11.5.1 Subgaussian Random Variables

Recall the definitions of the log-MGF $\psi(t)$ and its Legendre-Fenchel dual $\psi^*(\lambda)$ from §10.3.2.

Definition 11.5. A random variable V with mean 0 is *subgaussian with parameter σ* if its log-MGF $\psi(t)$ satisfies

$$\psi(t) \leq \frac{\sigma^2 t^2}{2}.$$

for all $t \geq 0$. It is *subgaussian with parameter σ up to t_0* if the above inequality holds for all $|t| \leq t_0$.

In other words, the log-MGF of a subgaussian r.v. is bounded above by that of a Gaussian! At this point, it's useful to recall a fact we asked as an exercise in §10.3.2:

Fact 11.6. If $\psi_1(t) \geq \psi_2(t)$ for all $t \geq 0$, then $\psi_1^*(\lambda) \leq \psi_2^*(\lambda)$ for all λ .

Using this, the dual function of a subgaussian random variable with parameter σ is bounded *below* by that of a Gaussian $N(0, \sigma^2)$, which means we have a tighter upper tail bound! Indeed, combining with (10.22), we immediately get:

Theorem 11.7 (Subgaussian Tail Bounds). *If V is zero-mean and subgaussian with parameter σ , then*

$$\Pr[V \geq \lambda] \leq e^{-\lambda^2/(2\sigma^2)}.$$

Most tail bounds you will prove using the subgaussian perspective will come down to showing that some random variable is subgaussian with parameter σ , whereupon you can use Theorem 11.7. Given that you will often reason about sums of subgaussians, you may use the next fact, which is an analog of Proposition 11.3.

Lemma 11.8. *If V_1, V_2, \dots are independent, zero-mean and σ_i -subgaussian, and x_1, x_2, \dots are reals, then $V = \sum_i x_i V_i$ is $\sqrt{\sum_i x_i^2 \sigma_i^2}$ -subgaussian.*

Proof. Using independence and the definition of subgaussian-ness:

$$\mathbb{E}[e^{tV}] = \mathbb{E}[e^{t \sum_i x_i V_i}] = \prod_i \mathbb{E}[e^{t x_i V_i}] \leq \prod_i e^{(t x_i)^2 \sigma_i^2 / 2}.$$

Finally taking logarithms, $\psi_V(t) = \sum_i \psi_{V_i}(t x_i) \leq \sum_i \frac{t^2 x_i^2 \sigma_i^2}{2}$. \square

11.6 A Proof of Lemma 11.2 using Subgaussian r.v.s

Suppose we choose each M_{ij} to be an independent copy of a subgaussian r.v. with zero mean and unit variance, and let $A(x) = \frac{1}{\sqrt{k}} Mx$ again? We want to show that

$$Z := \|A(x)\|^2 = \frac{1}{k} \sum_{i=1}^k (Mx)_i^2 \tag{11.9}$$

has mean $\|x\|^2$, and is concentrated sharply around that value. Conveniently, we had only used the mean and variance of the entries of M in proving Lemma 11.4, so we can still infer that

$$\mathbb{E}[Z] = \mathbb{E}[\|A(x)\|^2] = \|x\|^2 = 1.$$

It just remains to show the concentration.

11.6.1 Sums of Squares of Subgaussians

To add in. Until then see the explanation in Matousek's paper "On Variants of the Johnson-Lindenstrauss Lemma".

11.6.2 Relating Subgaussian to Gaussians

If you have done the proof for the Gaussian case, and just want to extend the JL Lemma to other subgaussian random variables, you need not do all the work in §11.6.1. Instead you can relate subgaussian concentration to good old Gaussian concentration.

Indeed, the direct proof from §11.4 showed the $(Mx)_i$ s were themselves Gaussian with variance $\|x\|^2$. Since the Rademachers are 1-subgaussian, Lemma 11.8 shows that $(Mx)_i$ is subgaussian with parameter $\|x\|^2$. Next, we need to consider Z , which is the average of squares of k independent $(Mx)_i$ s. The following lemma shows that the MGF of squares of *symmetric* σ -subgaussians are bounded above by the corresponding Gaussians with variance σ^2 .

Lemma 11.9. *If V is symmetric mean-zero σ -subgaussian r.v., and $W \sim N(0, \sigma^2)$, then $\mathbb{E}[e^{tV^2}] \leq \mathbb{E}[e^{tW^2}]$ for $t > 0$.*

Proof. Using the calculation in (10.20) in the "backwards" direction

$$\mathbb{E}_V[e^{tV^2}] = \mathbb{E}_{V,W}[e^{\sqrt{2t}(|V|/\sigma)W}].$$

(Note that we've just introduced W into the mix, without any provocation!) Since W is also symmetric, we get $|V|W = V|W|$. Hence, rewriting

$$\mathbb{E}_{V,W}[e^{\sqrt{2t}(|V|/\sigma)W}] = \mathbb{E}_W[\mathbb{E}_V[e^{(\sqrt{2t}|W|/\sigma)V}]],$$

we can use the σ -subgaussian behavior of V in the inner expectation to get an upper bound of

$$\mathbb{E}_W[e^{\sigma^2(\sqrt{2t}|W|/2)^2/2}] = \mathbb{E}_W[e^{tW^2}]. \quad \square$$

Excellent. Now the tail bound for sums of squares of symmetric mean-zero σ -subgaussians follows from that of Gaussians. Hence we get the same tail bounds as in §11.4.1, and hence that the Rademacher

An r.v. X is *symmetric* if it is distributed the same as $R|X|$, where R is an independent Rademacher.

matrix also has the distributional JL property, while using far fewer random bits!

In general one can use other σ -subgaussian distributions to fill the matrix M —using σ different than 1 may require us to rework the proof from §11.4.1 since the linear terms in (11.6) don’t cancel any more, see works by [Indyk and Naor](#) or [Matousek](#) for details.

[Indyk and Naor \(2008\)](#)
[Matoušek \(2008\)](#)

11.6.3 The Fast JL Transform

A different direction to consider is getting fast algorithms for the JL Lemma: Do we really need to plug in non-zero values into every entry of the matrix A ? What if most of A is filled with zeroes? The first problem is that if x is a very sparse vector, then Ax might be zero with high probability? Achlioptas showed that having a random two-thirds of the entries of A being zero still works fine: [Nir Ailon and Bernard Chazelle](#) showed that if you first hit x with a suitable matrix P which caused Px to be “well-spread-out” whp, and then $\|APx\| \approx \|x\|$ would still hold for a much sparser A . Moreover, this P requires much less randomness, and furthermore, the computations can be done faster too! There has been much work on fast and sparse versions of JL: see, e.g., this [paper from SOSA 2018](#) by Michael Cohen, T.S. Jayram, and Jelani Nelson. Jelani Nelson also has some [notes](#) on the Fast JL Transform.

[Ailon and Chazelle](#)

[Cohen, Jayram, and Nelson \(2018\)](#)

11.7 Optional: Compressive Sensing

[To rewrite](#). In an attempt to build a better machine to take MRI scans, we decrease the number of sensors. Then, instead of the signal x we intended to obtain from the machine, we only have a small number of measurements of this signal. Can we hope to recover x from the measurements we made if we make sparsity assumptions on x ? We use the term s -sparse signal for a vector with at most s nonzero entries, i.e., with $|\text{supp}(x)| \leq s$.

Formally, x is a n -dimensional s -sparse vector, and a measurement of x with respect to a vector a is a real number given by $\langle a, x \rangle$. If we ask k questions, this gives us a $k \times n$ *sensing matrix* A (whose rows are the measurements), and a k -dimensional vector b of results. We want to reconstruct x with s nonzero entries satisfying $Ax = b$. This is often written as

$$\min \left\{ \|x\|_0 \mid Ax = b \right\}. \quad (11.10)$$

It is common to use the notation $\|x\|_0 := |\text{supp}(x)|$, even though this is not a norm.

11.7.1 Sparse Recovery: A First Attempt

What properties would we like from our sensing matrix A ? The first would be some form of consistency: that the problem should be solvable.

Definition 11.10 (Kruskal Rank). An $m \times n$ matrix A has Kruskal rank r if every subset of r of its columns are linearly independent.

Lemma 11.11 (Unique Decoding). *If A has Kruskal rank $\geq 2s$, then for any b we have $Ax = b$ for at most one s -sparse x .*

Proof. Suppose $Ax = Ax'$ for two s -sparse vectors x, x' . Then $A(x - x') = 0$ for the $2s$ -sparse vector $x - x'$. The Kruskal rank being $2s$ means this vector $x - x' = 0$, and hence $x = x'$. \square

So we can just find some sensing matrix with large Kruskal rank [Give examples here](#) and ensure our results will be unique. The next question is: how fast can we find x ? (We should also be worried about noise in the measurements.) A generic construction of matrices with large Kruskal rank may not give us efficient solutions to (11.10). Indeed, it turns out that the problem as formulated is NP-hard, assuming A and b are contrived by an adversary.

Of course, asking to solve (11.10) for general A, b is a more difficult problem than we need to solve. In our setting, we can choose A as we like and then are given $b = Ax$, so we can ask whether there are matrices A for which this decoding process is indeed efficient. This is precisely what we do next.

11.7.2 The Basis Pursuit Algorithm

Consider the following similar looking problem called the *basis pursuit* (BP) problem:

$$\min \left\{ \|x\|_1 \mid Ax = b \right\}. \quad (11.11)$$

This problem can be formulated as a linear program as follows, and hence can be efficiently solved. Introduce n new variables y_1, y_2, \dots, y_n under the constraints

$$\min \left\{ \sum_i y_i \mid Ax = b, -y_i \leq x_i \leq y_i \right\}.$$

Definition 11.12. We call a matrix A as *BP-exact for sparsity s* if for all vectors b such that the non-convex program (11.10) has a unique solution x^* , this vector x^* is also the unique optimal solution to the basis pursuit LP (11.11).

In other words, we want a matrix A for which the two programs return the same optimal solution. But do BP-exact matrices exist? If so, how do we efficiently construct them? Our next ingredient will be crucial to show their existence and construction.

Definition 11.13 (Restricted Isometry Property (RIP)). A matrix A is (t, ε) -RIP if for all unit vectors x with $\|x\|_0 \leq t$, we have

$$\|Ax\|_2^2 \in [1 \pm \varepsilon].$$

Lemma 11.14 (RIP \implies BP-exact). *If a matrix A is $(3s, \varepsilon)$ -RIP for some $\varepsilon \leq 1/9$, then A is BP-exact for sparsity s .*

Proof. Suppose x^* is the unique solution to (11.10) and x the solution to (11.11), so that

$$\|x\|_1 \leq \|x^*\|_1. \quad (11.12)$$

Suppose $x - x^* = \Delta \neq 0$; hence $A\Delta = A(x - x^*) = 0$. If we could somehow show that $\text{supp}(\Delta) \leq 3s$, then using the RIP property for A , we would get

$$0 = \|A\Delta\|_2 \geq (1 - \varepsilon)\|\Delta\|_2 > 0,$$

a contradiction. But of course, Δ could have large support, so we need to work harder. The actual proof breaks up Δ into small pieces (so that the RIP matrix A maintains their length), and argues that there is one large piece that the other pieces cannot cancel out.

Let $S := \text{supp}(x^*)$ be the support of x^* , and \bar{S} be the remaining coordinates. Let's sort these coordinates in decreasing order of their *absolute value*, and group them into buckets of $2s$ consecutive coordinates. Call these buckets B_1, B_2, \dots

Claim 11.15. $\sum_{j \geq 2} \|\Delta_{B_j}\|_2 \leq \|\Delta_S\|_2 / \sqrt{2}$.

Before we prove the claim, let's see how to use it. The claim says that total Euclidean length of the vectors $\{v_{B_j}\}_{j \geq 2}$ is a constant factor smaller than that of $v_{S \cup B_1}$. So even after the near-isometric mapping A , the lengths of the former would not be able to cancel the length of the latter. Formally:

$$\begin{aligned} 0 &= \|A\Delta\|_2 \geq \|A\Delta_{S \cup B_1}\|_2 - \sum_{j \geq 2} \|A\Delta_{B_j}\|_2 \\ &\geq (1 - \varepsilon) \|\Delta_{S \cup B_1}\|_2 - (1 + \varepsilon) \sum_{j \geq 2} \|\Delta_{B_j}\|_2 \\ &\geq (1 - \varepsilon) \|\Delta_S\|_2 - \frac{1 + \varepsilon}{\sqrt{2}} \|\Delta_S\|_2, \end{aligned}$$

where the first step uses the triangle inequality for norms, the second uses that each $\Delta_{S \cup B_1}$ and Δ_{B_j} are $3s$ -sparse, and the last step uses

For vector $v \in \mathbb{R}^n$ and subset $T \subseteq [n]$, define vector $v_T \in \mathbb{R}^n$ which agrees with v on the coordinates in T , and which has zeroes elsewhere.

$\|\Delta_{S \cup B_1}\|_2 \geq \|\Delta_S\|_2$ and also Claim 11.15. Finally, since $\varepsilon \leq 1/9$, we have $1 - \varepsilon > \frac{1+\varepsilon}{\sqrt{2}}$, so the only remaining possibility is that $\Delta_S = 0$. The next claim implies that $\Delta_S = 0$ implies that $\Delta = 0$, giving a contradiction and hence the proof of Lemma 11.14.

Claim 11.16. $\|\Delta_S\|_1 \geq \|\Delta_{\bar{S}}\|_1$.

Proof. We finally use that $x = x^* + \Delta$ is the optimizer for the LP, which means

$$\begin{aligned}\|x^*\|_1 &> \|x^* + \Delta\|_1 = \|x_S^* + \Delta_S\|_1 + \|\Delta_{\bar{S}}\|_1 \\ &\geq \|x_S^*\|_1 - \|\Delta_S\|_1 + \|\Delta_{\bar{S}}\|_1.\end{aligned}$$

(The last step uses the triangle inequality.) Since $\|x^*\|_1 = \|x_S^*\|_1$, we get Claim 11.16. \square

The final piece of the argument is to prove Claim 11.15:

Proof of Claim 11.15. Take any bucket B_j for $j \geq 2$. Each entry of Δ in this bucket is smaller than the smallest entry of B_{j-1} , and hence smaller than the average entry of B_{j-1} . And there are $2s$ entries in this bucket B_j , so the Euclidean length of the bucket is

$$\|\Delta_{B_j}\|_2 \leq \sqrt{2s} \cdot \frac{\|\Delta_{B_{j-1}}\|_1}{2s} = \frac{\|\Delta_{B_{j-1}}\|_1}{\sqrt{2s}}.$$

Summing this over all $j \geq 2$, we get

$$\sum_{j \geq 2} \|\Delta_{B_j}\|_2 \leq \sum_{j \geq 2} \frac{\|\Delta_{B_{j-1}}\|_1}{\sqrt{2s}} = \frac{\|\Delta_{\bar{S}}\|_1}{\sqrt{2s}}.$$

Now $\|\Delta_{\bar{S}}\|_1 \leq \|\Delta_S\|_1$ by Claim 11.16. And finally, since the support of Δ_S is of size s , we can bound its ℓ_1 length by \sqrt{s} times its ℓ_2 length, finishing the claim. (Since we wanted that factor of $\sqrt{2}$ in the denominator, we made the buckets slightly larger than the size of S .) \square

Exercise: for any vector $v \in \mathbb{R}^d$, show that $\|v\|_1 \leq \sqrt{\text{supp}(v)} \cdot \|v\|_2$.

This completes the proof for Lemma 11.14. \square

Finally, how do we construct RIP matrices? Call a distribution \mathcal{D} over $k \times n$ matrices a *distributional JL family* if Lemma 11.2 is true when A is drawn from \mathcal{D} . The following theorem was proved by David Donoho, and by Emmanuel Candes and Terry Tao, and by Mark Rudelson and Roman Vershynin. (The connection of their construction to the distributional JL was made explicit by Baraniuk et al.)

Theorem 11.17 (JL \implies RIP). *If we pick $A \in \mathbb{R}^{k \times n}$ from a distributional JL family with $k \geq \Omega(s \log n/s)$, then with high probability A is BP-exact.*

Proof. The proof is simple, but uses some fairly general ideas worth emphasizing. First, focus on some s -dimensional subspace of \mathbb{R}^n (obtained by restricting to some subset of coordinates). For notational simplicity, we just identify this subspace with \mathbb{R}^s .

1. For $\delta = \varepsilon/3$, pick an δ -net N of the sphere S^{s-1} (under Euclidean distances). This can be done by a greedy algorithm: if some point x does not satisfy the covering property at any time, it can be added to the net. We claim the size of the net is $|N| := (4/\delta)^s$. Indeed, define balls of radius $\delta/2$ around the points in N ; these are disjoint by the packing property of nets, and are all contained in a ball of radius $1 + \delta$ around the origin. Since the volume of balls of radius r scales as r^s , we have

$$|N| \leq \left(\frac{1+\delta}{\delta/2}\right)^s = (4/\delta)^s.$$

2. If A is an δ -isometry on the δ -net $N \subseteq S^{s-1}$, we claim it is a 3δ -isometry on all of S^{s-1} . Indeed, consider the point x that achieves the maximum stretch $\arg \max \{\|Ax\|_2 \mid x \in S^{s-1}\}$, and let this stretch be M . Let y be the closest point in N to x ; by the packing property $\|x - y\| \leq \delta$. Then $M = \|Ax\| \leq \|Ay\| + \|A(x - y)\| \leq (1 + \delta) + M\delta$. Rearranging, $M \leq \frac{1+\delta}{1-\delta} \leq (1 + 3\delta)$ for $\delta \leq 1/3$, say. For the contraction, consider any $x \in S^{s-1}$, with closest net point y . Then $\|Ax\| \geq \|Ay\| - \|A(x - y)\| \geq 1 - \delta - (1 + 3\delta)\delta \geq 1 - 3\delta$, again as long as $\delta \leq 1/3$.
3. By Lemma 11.2, the random matrix A with m rows is an δ -isometry on each point in the net N , except with probability $\exp(-c\delta^2 m)$ for some constant c .
4. Now apply the above argument to each of the $\binom{n}{s}$ subspaces obtained by restricting to some subset S of coordinates. By a union bound over all subsets S , and over all points in the net for that subspace, the matrix A is an 3δ -isometry on all points with support in S except with probability

$$\binom{n}{s} \cdot (4/\delta)^s \cdot \exp(-c\delta^2 m) \leq \exp(-\Theta(m)),$$

as long as m is $\Omega(s \log n/s)$. Since $\varepsilon = 3\delta$, we have the proof. \square

Given a metric space (X, d) , a δ -net is a subset $N \subseteq X$ such that (i) $d(x, y) \geq \delta$ for all $x, y \in N$, and (ii) for each $x \in X$ there exists $y \in N$ such that $d(x, y) \leq \delta$. The former is called the *packing property* and the latter the *covering property* of nets.

This presentation is based on [notes](#) by Jirka Matoušek. Also see Chapter 4 of [Ankur Moitra's book](#) for more on compressed sensing, sparse recovery and basis pursuit.

11.8 Some Facts about Balls in High-Dimensional Spaces

Consider the unit ball $\mathbb{B}_d := \{x \in \mathbb{R}^d \mid \|x\|_2 \leq 1\}$. Here are two facts, whose proofs we sketch. These sketches can be made formal (since the approximations are almost the truth), but perhaps the style of arguments are more illuminating.

Theorem 11.18 (Heavy Shells). *At least $1 - \varepsilon$ of the mass of the unit ball in \mathbb{R}^d lies within a $\Theta(\frac{\log 1/\varepsilon}{d})$ -width shell next to the surface.*

Proof. (Sketch) The volume of a radius- r ball in \mathbb{R}^d goes as r^d , so the fraction of the volume *not* in the shell of width w is $(1-w)^d \approx e^{-wd}$, which is ε when $w \approx \frac{\log 1/\varepsilon}{d}$. \square

Given any hyperplane $H = \{x \in \mathbb{R}^d \mid a \cdot x = b\}$ where $\|a\| = 1$, the width- w slab around it is $K = \{x \in \mathbb{R}^d \mid b-w \leq a \cdot x \leq b+w\}$.

Theorem 11.19 (Heavy Slabs). *At least $(1 - \varepsilon)$ of the mass of the unit ball in \mathbb{R}^d lies within $\Theta(1/\sqrt{d})$ slab around any hyperplane that passes through the origin.*

Proof. (Sketch) By spherical symmetry we can consider the hyperplane $\{x_1 = 0\}$. The volume of the ball within $\{-w \leq x_1 \leq w\}$ is at least

$$\int_{y=0}^w (\sqrt{1-y^2})^{d-1} dy \approx \int_{y=0}^w e^{-y^2 \cdot \frac{d-1}{2}} dy.$$

If we define $\sigma^2 = \frac{1}{d-1}$, this is

$$\int_{y=0}^w e^{-\frac{y^2}{2\sigma^2}} dy \approx \Pr[G \leq w],$$

where $G \sim N(0, \sigma^2)$. But we know that $\Pr[G \geq w] \leq e^{-w^2/2\sigma^2}$ by our generic Chernoff bound for Gaussians (10.21). So setting that tail probability to be ε gives

$$w \approx \sqrt{2\sigma^2 \log(1/\varepsilon)} = O\left(\sqrt{\frac{\log(1/\varepsilon)}{d}}\right). \quad \square$$

This may seem quite counter-intuitive: that 99% of the volume of the sphere is within $O(1/d)$ of the surface, yet 99% is within $O(1/\sqrt{d})$ of *any* central slab! This challenges our notion of the ball “looking like” the smooth circular object, and more like a very spiky sea-urchin. Finally, a last observation:

Corollary 11.20 (Near-orthogonality). *Two random vectors from the surface of the unit ball in \mathbb{R}^d (i.e., from the sphere S^{d-1}) are nearly orthogonal with high probability. In particular, their dot-product is smaller than $O(\sqrt{\frac{\log(1/\varepsilon)}{d}})$ with probability $1 - \varepsilon$.*

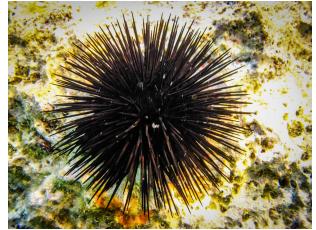


Figure 11.1: Sea Urchin (from uncommoncaribbean.com)

Proof. Fix one of the vectors u . Then for dot-product $|u \cdot v|$ to be at most ε , the other vector v must fall in the slab of width ε around the hyperplane $\{x \cdot u = 0\}$. Now Theorem 11.19 completes the argument. \square

This means that if we pick n random vectors in \mathbb{R}^d , and set $\varepsilon = 1/n^2$, a union bound gives that all have dot-product $O(\sqrt{\frac{\log n}{d}})$. Setting this dot-product to ε gives us $n = \exp(\varepsilon^2 d)$ unit vectors with mutual dot-products at most ε , exactly as in the calculation at the beginning of the chapter.