Lecture Notes for STAT240 (Robust and Nonparametric Statistics)

Jacob Steinhardt

Last updated: January 17, 2021

[Lecture 1]

1 What is this course about?

Consider the process of building a statistical or machine learning model. We typically first collect training data, then fit a model to that data, and finally use the model to make predictions on new test data.

In theory and in practice, we generally rely on the train and test data coming from the same distribution, or at least being closely related in some way. However, there are several ways this could fail to be the case:

1. The data collection process itself could be noisy and thus not reflect the actual underlying signal we wish to learn. For instance, there could be human error in labelling or annotation, or measurement error due to imperfect sensors.

2. There could be distributional shift, due to changes in the world over time or because we seek to deploy the model in some new situation (a language model trained on news articles but deployed on twitter). There might also be noise e.g. due to a sensor failing.

Robustness concerns what we should do when the train and test distribution are not the same, for any of the reasons above. There are two underlying perspectives influencing the choice of material in this course. First, we are generally interested in *worst-case* rather than average-case robustness. For instance, when handling data collection errors we will avoid modeling the errors as random noise and instead build procedures that are robust to any errors within some allowed family. We prefer this because average-case robustness requires the errors to satisfy a single, specific distribution for robustness guarantees to be meaningful, while a goal of robustness is to handle unanticipated situations that are difficult to model precisely in advance.

Second, we will study robustness in *high-dimensional* settings. Many natural approaches to robustness that work in low dimensions fail in high dimensions. For instance, the median is a robust estimate of the mean in one dimension, but the per-coordinate median is a poor robust estimator when the dimension is large (its error grows as $\sqrt{d}$ in d dimensions). We will see that more sophisticated estimators can substantially improve on this first attempt.

Complementary to robustness is the idea of *model mis-specification*. When the true distribution p^* lies within our model family, many robustness issues are less severe: we can rely on the model to extrapolate to new settings, and we can often get well-calibrated uncertainty estimates. This motivates the second focus of the course, *nonparametric modeling*, where we consider broad function classes (e.g. all smooth functions) that are more likely to be correctly specified. Another connection between nonparametrics and robustness is that we often want robust methods to work for any distribution within some large, infinite-dimensional class.

Overarching framework. Most robustness questions can be cast in the following way: We let p^* denote the true test distribution we wish to estimate, and assume that training data $X_1, \ldots, X_n$ is sampled i.i.d. from some distribution $\tilde{p}$ such that $D(\tilde{p}, p^*) \leq \epsilon$ according to some discrepancy D. We also assume that $p^* \in \mathcal{G}$, which encodes the distributional assumptions we make (e.g. that p^* has bounded moments or tails, which is typically necessary for robust estimation to be possible). We benchmark an estimator $\hat{\theta}(X_1, \ldots, X_n)$ according to some cost $L(p^*, \hat{\theta})$ (the test error). The diagram in Figure 1 illustrates this.

This framework captures both of the examples discused at the beginning. However, it will be profitable to think about each case separately due to different emphases:

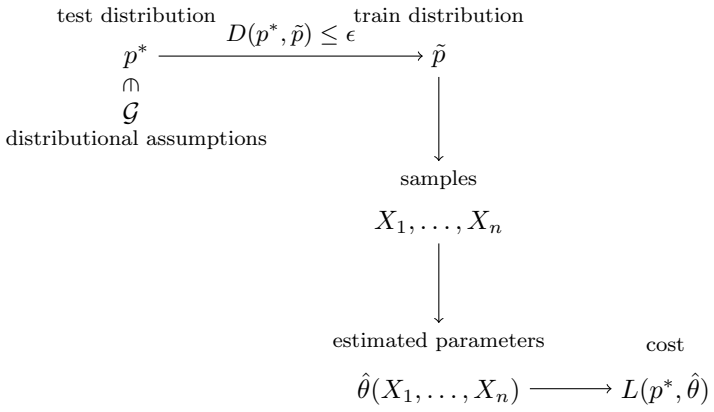

Figure 1: Framework for studying robust learning.

1. For corrupted training data, we think of $\tilde{p}$ as being corrupted and p^* as being nice.
2. For distributional shift, we think of $\tilde{p}$ and p^* as both being nice (but different).

Additionally, since both $\tilde{p}$ and p^* are nice for distributional shift, we should have greater ambitions and seek to handle larger differences between train and test than in the corruption cases.

Training robustness. Designing robust estimators for training corruptions usually involves reasoning about what the real data "might have" looked like. This could involve operations such as removing outliers, smoothing points away from extremes, etc. Unfortunately, many intuitive algorithms in low dimensions achieve essentially trivial bounds in high dimensions. We will show how to achieve more meaningful bounds, focusing on three aspects:

1. good dependence of the error on the dimension,
2. good finite-sample bounds,
3. computational tractability.

Each of these aspects turns out to require new machinery and we will devote roughly equal space to each.

Distributional shift. For distributional shift, we often seek invariant features or structure that can transfer information from the train to test distributions. We can also counteract distribution shift by training on more diverse data. Finally, we can use model uncertainty to infer out-of-distribution error, but these inferences can be fraught if the model is mis-specified.

Table 1: Comparison of different robust settings.

Train robustness	Distributional shift
p^* nice	p^* and $\tilde{p}$ both nice
undo corruptions	invariant features
	diverse training data

Nonparametric modeling. This brings us to nonparametric methods. The simplest way to be nonparametric is through the model–using a rich class such as smooth functions, or neural networks. This mitigates model mis-specification but raises new statistical challenges. Familiar phenomena such as the central limit theorem cease to hold, and error rates are instead governed by the eigenvalue spectrum of an infinite-dimensional kernel. Aside from the model, we can also be nonparametric at inference time; an example is the bootstrap, which constructs confidence intervals that are more robust to model mis-specification than classical parametric tests.

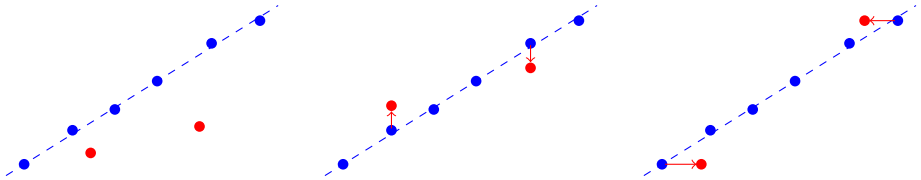

Figure 2: Possible corruptions to be robust to. Left: data contains outliers. Middle: outputs are perturbed (process noise); Right: inputs are perturbed (measurement error).

Summary. We will cover each of training time robustness, distribution shift, and nonparametric modeling, drawing connections as we go. The first third will focus on training time, also building up the statistical machinery needed to prove generalization bounds. Then, we will shift focus to model mis-specification, and to nonparametric methods as a remedy. These including kernel regression, the bootstrap, and neural networks, and we will both see how they mitigate model mis-specification and how to analyze their generalization performance. Finally, we will turn to distribution shift, and see that nonparametric models are often robust to distribution shift when trained on the right data. Training robustness will receive a largely theoretical treatment, while for nonparametrics and distribution shift we will see a mix of theoretical and empirical results.

2 Training Time Robustness

We will start our investigation with training time robustness. As in Figure 1, we observe samples $X_1, \dots, X_n$ from a corrupted training distribution $\tilde{p}$, whose relationship to the true (test) distribution is controlled by the constraint $D(\tilde{p}, p^*) \le \epsilon$. We additionally constrain $p^* \in \mathcal{G}$, which encodes our distributional assumptions.

Note that this setting corresponds to an *oblivious* adversary that can only apply corruptions at the population level (changing p^* to $\tilde{p}$); we can also consider a more powerful *adaptive* adversary that can apply corruptions to the samples themselves. Such an adversary is called adaptive because it is allowed to adapt to the random draw of the samples points $X_1, \dots, X_n$. Formally defining adaptive adversaries is somewhat technical and we defer this to later.

Figure 2 illustrates several ways in which a training distribution could be corrupted. In the left panel, an ϵ fraction of real points have been replaced by outliers. This can be modeled by the discrepancy $D(p, q) = \mathsf{TV}(p, q)$, where TV is the *total variation distance*:

$$\mathsf{TV}(p, q) \stackrel{\text{def}}{=} \sup\{|p(E) - q(E)| \mid E \text{ is a measurable event}\}. \tag{1}$$

If p and q both have densities then an equivalent characterization is $\mathsf{TV}(p, q) = \frac{1}{2} \int |p(x) - q(x)| dx$.

In the middle and right panels of Figure 2, either the inputs or outputs have been moved slightly. Both operations can be modeled using *Wasserstein distances* (also called earthmover distances), which we will discuss later. For now, however, we will focus on the case of handling outliers. Thus for the next several sections our discrepancy will be the total variation distance $D = \mathsf{TV}$.

2.1 Robustness to Outliers in 1 Dimension

First consider mean estimation in one dimension: we observe n data points $x_1, \dots, x_n \in \mathbb{R}$ drawn from $\tilde{p}$, and our goal is to estimate the mean $\mu = \mathbb{E}_{x \sim p^*}[x]$ of p^*. Accordingly our loss is $L(p^*, \theta) = |\theta - \mu(p^*)|$.

The following histogram illustrates a possible dataset, where the height of each bar represents the number of points with a given value:

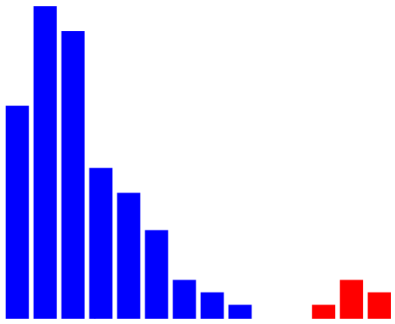

Are the red points outliers? Or part of the real data? Depending on the conclusion, the estimated mean could vary substantially. Without further assumptions on the data-generating distribution p^*, we cannot rule out either case. This brings us to an important principle:

With no assumptions on the distribution p^*, robust estimation is impossible.

In particular, we must make assumptions that are strong enough to reject sufficiently extreme points as outliers, or else even a small fraction of such points can dominate the estimate of the mean. For simplicity, here and in the next several sections we will assume that we directly observe the training distribution $\tilde{p}$ rather than samples $x_{1:n}$ from $\tilde{p}$. This allows us to avoid analyzing finite-sample concentration, which requires introducing additional technical tools that we will turn to in Section 2.5.

Assumption: bounded variance. One possible assumption is that p^* has bounded variance: $\mathbb{E}_{x\sim p^*}[(x-\mu)^2] \leq \sigma^2$ for some parameter σ. We take $\mathcal{G} = \mathcal{G}_{\mathsf{cov}}(\sigma)$ to be the set of distributions satisfying this constraint.

Under this assumption, we can estimate μ to within error $\mathcal{O}(\sigma\sqrt{\epsilon})$ under TV-perturbations of size ϵ. Indeed, consider the following procedure:

Algorithm 1 `TrimmedMean`

1: Remove the upper and lower (2ϵ)-quantiles from $\tilde{p}$ (so 4ϵ mass is removed in total).
2: Let $\tilde{p}_{2\epsilon}$ denote the new distribution after re-normalizing, and return the mean of $\tilde{p}_{2\epsilon}$.

To analyze Algorithm 1, we will make use of a strengthened version of Chebyshev's inequality, which we recall here (see Section B.1 for a proof):

Lemma 2.1 (Chebyshev inequality)**.** *Suppose that p has mean μ and variance σ^2. Then, $\mathbb{P}_{X\sim p}[X \geq \mu + \sigma/\sqrt{\delta}] \leq \delta$. Moreover, if E is any event with probability at least δ, then $|\mathbb{E}_{X\sim p}[X \mid E] - \mu| \leq \sigma\sqrt{\frac{2(1-\delta)}{\delta}}$.*

The first part, which is the standard Chebyshev inequality, says that it is unlikely for a point to be more than a few standard deviations away from μ. The second part says that any large population of points must have a mean close to μ. This second property, which is called *resilience*, is central to robust estimation, and will be studied in more detail in Section 2.4.

With Lemma 2.1 in hand, we can prove the following fact about Algorithm 1:

Proposition 2.2. *Assume that $\mathsf{TV}(\tilde{p}, p^*) \leq \epsilon \leq \frac{1}{8}$. Then the output $\hat{\mu}$ of Algorithm 1 satisfies $|\hat{\mu} - \mu| \leq 9\sigma\sqrt{\epsilon}$.*

Proof. If $\mathsf{TV}(\tilde{p}, p^*) \leq \epsilon$, then we can get from p^* to $\tilde{p}$ by adding an ϵ-fraction of points (outliers) and deleting an ϵ-fraction of the original points.

First note that all outliers that exceed the ϵ-quantile of p^* are removed by Algorithm 1. Therefore, all non-removed outliers lie within $\frac{\sigma}{\sqrt{\epsilon}}$ of the mean μ by Chebyshev's inequality.

Second, we and the adversary together remove at most a 5ϵ-fraction of the mass in p^*. Applying Lemma 2.1 with $\delta = 1 - 5\epsilon$, the mean of the remaining good points lies within $\sigma\sqrt{\frac{10\epsilon}{1-5\epsilon}}$ of μ.

Now let ϵ' be the fraction of remaining points which are bad, and note that $\epsilon' \leq \frac{\epsilon}{1-4\epsilon}$. The mean of all the remaining points differs from μ by at most $\epsilon' \cdot \sigma\sqrt{\frac{1}{\epsilon}} + (1-\epsilon') \cdot \sigma\sqrt{\frac{10\epsilon}{1-5\epsilon}}$, which is at most $(1+\sqrt{10})\frac{\sqrt{\epsilon}}{1-4\epsilon}\sigma$. This is in turn at most $9\sigma\sqrt{\epsilon}$ assuming that $\epsilon \leq \frac{1}{8}$. □

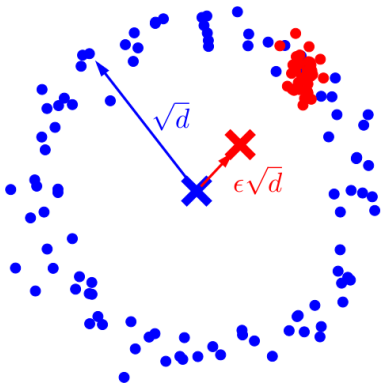

Figure 3: The outliers can lie at distance $\sqrt{d}$ without being detected, skewing the mean by $\epsilon\sqrt{d}$.

Optimality. The $\mathcal{O}(\sigma\sqrt{\epsilon})$ dependence is optimal, because the adversary can themselves apply the same trimming procedure we do, and in general this will shift the mean of a bounded covariance distribution by $\mathcal{O}(\sigma\sqrt{\epsilon})$ while keeping the covariance bounded.

Alternate assumptions. The key fact driving the proof of Proposition 2.2 is that any $(1-\epsilon)$-fraction of the good points has mean at most $\mathcal{O}(\sigma\sqrt{\epsilon})$ away from the true mean due to Chebyshev's inequality (Lemma 2.1), which makes use of the bound σ^2 on the variance. Any other bound on the deviation from the mean would yield an analogous result. For instance, if p^* has bounded kth moment, then the $\mathcal{O}(\sigma\sqrt{\epsilon})$ in Lemma 2.1 can be improved to $\mathcal{O}(\sigma_k\epsilon^{1-1/k})$, where $(\sigma_k)^k$ is a bound on the kth moment; in this case Algorithm 1 will estimate μ with a correspondingly improved error of $\mathcal{O}(\sigma_k\epsilon^{1-1/k})$.

2.2 Problems in High Dimensions

In the previous section, we saw how to robustly estimating the mean of a 1-dimensional dataset assuming the true data had bounded variance. Our estimator worked by removing data points that are too far away from the mean, and then returning the mean of the remaining points.

It is tempting to apply this same idea in higher dimensions—for instance, removing points that are far away from the mean in ℓ_2-distance. Unfortunately, this incurs large error in high dimensions.

To see why, consider the following simplified example. The distribution p^* over the true data is an isotropic Gaussian $\mathcal{N}(\mu, I)$, with unknown mean μ and independent variance 1 in every coordinate. In this case, the typical distance $\|x_i - \mu\|_2$ of a sample x_i from the mean μ is roughly $\sqrt{d}$, since there are d coordinates and x_i differs from μ by roughly 1 in every coordinate. (In fact, $\|x_i - \mu\|_2$ can be shown to concentrate around $\sqrt{d}$ with high probability.) This means that the outliers can lie at a distance $\sqrt{d}$ from μ without being detected, thus shifting the mean by $\Theta(\epsilon\sqrt{d})$; Figure 3 depicts this. Therefore, filtering based on ℓ_2 distance incurs an error of at least $\epsilon\sqrt{d}$. This dimension-dependent $\sqrt{d}$ factor often renders bounds meaningless.

In fact, the situation is even worse; not only are the bad points no further from the mean than the good points in ℓ_2-distance, they actually have the same probability density under the true data-generating distribution $\mathcal{N}(\mu, I)$. There is thus no procedure that measures each point in isolation and can avoid the $\sqrt{d}$ factor in the error.

This leads us to an important take-away: *In high dimensions, outliers can substantially perturb the mean while individually looking innocuous.* To handle this, we will instead need to analyze entire populations of outliers at once. In the next section we will do this using *minimum distance functionals*, which will allow us to avoid the dimension-dependent error.

[Lecture 2]

2.3 Minimum Distance Functionals

In the previous section we saw that simple approaches to handling outliers in high-dimensional data, such as the trimmed mean, incur a $\sqrt{d}$ error. We will avoid this error using *minimum distance functionals*, an idea which seems to have first appeared in Donoho and Liu (1988).

Definition 2.3 (Minimum distance functional)**.** For a family $\mathcal{G}$ and discrepancy D, the minimum distance functional is

$$\hat{\theta}(\tilde{p}) = \theta^*(q) = \arg\min_{\theta} L(q, \theta), \text{ where } q = \arg\min_{q \in \mathcal{G}} D(q, \tilde{p}). \tag{2}$$

In other words, $\hat{\theta}$ is the parameters obtained by first projecting $\tilde{p}$ onto $\mathcal{G}$ under D, and then outputting the optimal parameters for the resulting distribution.

An attractive property of the minimum-distance functional is that it does not depend on the perturbation level ϵ. More importantly, it satisfies the following cost bound in terms of the *modulus of continuity* of $\mathcal{G}$:

Proposition 2.4. *Suppose D is a pseudometric. Then the cost $L(p^*, \hat{\theta}(\tilde{p}))$ of the minimum distance functional is at most the maximum loss between any pair of distributions in $\mathcal{G}$ of distance at most 2ϵ:*

$$\mathfrak{m}(\mathcal{G}, 2\epsilon, D, L) \triangleq \sup_{p,q \in \mathcal{G}: D(p,q) \le 2\epsilon} L(p, \theta^*(q)). \tag{3}$$

The quantity $\mathfrak{m}$ is called the modulus of continuity because, if we think of $L(p, \theta^*(q))$ as a discrepancy between distributions, then $\mathfrak{m}$ is the constant of continuity between L and D when restricted to pairs of nearby distributions in $\mathcal{G}$.

Specialize again to the case $D = \mathsf{TV}$ and $L(p^*, \theta) = \|\theta - \mu(p^*)\|_2$ (here we allow p^* to be a distribution over $\mathbb{R}^d$ rather than just $\mathbb{R}$). Then the modulus is $\sup_{p,q \in \mathcal{G}: \mathsf{TV}(p,q) \le 2\epsilon} \|\mu(p) - \mu(q)\|_2$. As a concrete example, let $\mathcal{G}$ be the family of Gaussian distributions with unknown mean μ and identity covariance. For this family, the TV distance is essentially linear in the difference in mean:

Lemma 2.5. *Let $\mathcal{N}(\mu, I)$ denote a Gaussian distribution with mean μ and identity covariance. Then*

$$\min(u/2, 1)/\sqrt{2\pi} \le \mathsf{TV}(\mathcal{N}(\mu, I), \mathcal{N}(\mu', I)) \le \min(u/\sqrt{2\pi}, 1), \tag{4}$$

where $u = \|\mu - \mu'\|_2$.

Proof. By rotational and translational symmetry, it suffices to consider the case of one-dimensional Gaussians $\mathcal{N}(-u/2, 1)$ and $\mathcal{N}(u/2, 1)$. Then we have that

$$\mathsf{TV}(\mathcal{N}(-u/2, 1), \mathcal{N}(u/2, 1)) = \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\infty} |e^{-(t+u/2)^2/2} - e^{-(t-u/2)^2/2}| dt \tag{5}$$

$$\overset{(i)}{=} \frac{1}{\sqrt{2\pi}} \int_{-u/2}^{u/2} e^{-t^2/2} dt. \tag{6}$$

(The equality (i) is a couple lines of algebra, but is easiest to see by drawing a graph of the two Gaussians and cancelling out most of the probability mass.)

For the lower bound, note that $e^{-t^2/2} \ge \frac{1}{2}$ if $|t| \le 1$.

For the upper bound, similarly note that $e^{-t^2/2} \le 1$ for all $t \in \mathbb{R}$, and also that the entire integral must be at most 1 because it is the probability density of a Gaussian. □

Lemma 2.5 allows us to compute the modulus for Gaussians:

Corollary 2.6. *Let $\mathcal{G}_{\mathsf{gauss}}$ be the family of isotropic Gaussians, $D = \mathsf{TV}$, and L the difference in means as above. Then $\mathfrak{m}(\mathcal{G}_{\mathsf{gauss}}, \epsilon, D, L) \le 2\sqrt{2\pi}\epsilon$ whenever $\epsilon \le \frac{1}{2\sqrt{2\pi}}$.*

In particular, by Proposition 2.4 the minimum distance functional achieves error $\mathcal{O}(\epsilon)$ for Gaussian distributions when $\epsilon \le \frac{1}{2\sqrt{2\pi}}$. This improves substantially on the $\epsilon\sqrt{d}$ error of the trimmed mean estimator from Section 2.2. We have achieved our goal at least for Gaussians.

More general families. Taking $\mathcal{G}$ to be Gaussians is restrictive, as it assumes that p^* has a specific parametric form—counter to our goal of being robust! However, the modulus $\mathfrak{m}$ is bounded for much more general families. As one example, we can take the distributions with bounded covariance (compare to Proposition 2.2):

Lemma 2.7. *Let $\mathcal{G}_{\mathsf{cov}}(\sigma)$ be the family of distributions whose covariance matrix Σ satisfies $\Sigma \preceq \sigma^2 I$. Then $\mathfrak{m}(\mathcal{G}_{\mathsf{cov}}(\sigma), \epsilon) = \mathcal{O}(\sigma\sqrt{\epsilon})$.*

Proof. Let $p, q \in \mathcal{G}_{\mathsf{cov}}(\sigma)$ such that $\mathsf{TV}(p, q) \leq \epsilon$. This means that we can get from p to q by first deleting ϵ mass from p and then adding ϵ new points to end up at q. Put another way, there is a distribution r that can be reached from both p and q by deleting ϵ mass (and then renormalizing). In fact, this distribution is exactly

$$r = \frac{\min(p, q)}{1 - \mathsf{TV}(p, q)}. \tag{7}$$

Since r can be obtained from both p and q by deletions, we can make use of the following multi-dimensional analogue of Chebyshev's inequality (Lemma 2.1):

Lemma 2.8 (Chebyshev in $\mathbb{R}^d$)**.** *Suppose that p has mean μ and covariance Σ, where $\Sigma \preceq \sigma^2 I$. Then, if E is any event with probability at least δ, we have $\|\mathbb{E}_{X\sim p}[X \mid E] - \mu\|_2 \leq \sigma\sqrt{\frac{2(1-\delta)}{\delta}}$.*

As a consequence, we have $\|\mu(r) - \mu(p)\|_2 \leq \sigma\sqrt{2\epsilon/(1-\epsilon)}$ and $\|\mu(r) - \mu(q)\|_2 \leq \sigma\sqrt{2\epsilon/(1-\epsilon)}$ (since r can be obtained from either p or q by conditioning on an event of probability $1 - \epsilon$). By triangle inequality and assuming $\epsilon \leq \frac{1}{2}$, we have $\|\mu(p) - \mu(q)\|_2 \leq 4\sigma\sqrt{\epsilon}$, as claimed. □

As a consequence, the minimum distance functional robustly estimates the mean bounded covariance distributions with error $\mathcal{O}(\sigma\sqrt{\epsilon})$, generalizing the 1-dimensional bound obtained by the trimmed mean.

In Lemma 2.7, the two key properties we needed were:

- The *midpoint property* of TV distance (i.e., that there existed an r that was a deletion of p and q).
- The *bounded tails* guaranteed by Chebyshev's inequality.

If we replace bounded covariance distributions with any other family that has tails bounded in a similar way, then the minimum distance functional will similarly yield good bounds. A general family of distributions satisfying this property are *resilience distributions*, which we turn to next.

2.4 Resilience

Here we generalize Lemma 2.7 to prove that the modulus of continuity $\mathfrak{m}$ is bounded for a general family of distributions containins Gaussians, sub-Gaussians, bounded covariance distributions, and many others. The main observation is that in the proof of Lemma 2.7, all we needed was that the tails of distributions in $\mathcal{G}$ were bounded, in the sense that deleting an ϵ-fraction of the points could not substantially change the mean. This motivates the following definition:

Definition 2.9. A distribution p over $\mathbb{R}^d$ is said to be (ρ, ϵ)-resilient (with respect to some norm $\|\cdot\|$) if

$$\|\mathbb{E}_{X\sim p}[X \mid E] - \mathbb{E}_{X\sim p}[X]\| \leq \rho \text{ for all events } E \text{ with } p(E) \geq 1 - \epsilon. \tag{8}$$

We let $\mathcal{G}_{\mathsf{TV}}(\rho, \epsilon)$ denote the family of (ρ, ϵ)-resilient distributions.

We observe that $\mathcal{G}_{\mathsf{cov}}(\sigma) \subset \mathcal{G}_{\mathsf{TV}}(\sigma\sqrt{2\epsilon/(1-\epsilon)}, \epsilon)$ for all ϵ by Lemma 2.8; in other words, bounded covariance distributions are resilient. We can also show that $\mathcal{G}_{\mathsf{gauss}} \subset \mathcal{G}_{\mathsf{TV}}(2\epsilon\sqrt{\log(1/\epsilon)}, \epsilon)$, so Gaussians are resilient as well.

Resilient distributions always have bounded modulus:

Theorem 2.10. *The modulus of continuity $\mathfrak{m}(\mathcal{G}_{\mathsf{TV}}, 2\epsilon)$ satisfies the bound*

$$\mathfrak{m}(\mathcal{G}_{\mathsf{TV}}(\rho, \epsilon), 2\epsilon) \leq 2\rho \tag{9}$$

whenever $\epsilon < 1/2$.

Proof. As in Lemma 2.7, the key idea is that any two distributions p, q that are close in TV have a *midpoint* distribution $r = \frac{\min(p,q)}{1-\mathsf{TV}(p,q)}$ that is a deletion of both distributions). This midpoint distribution connects the two distributions, and it follows from the triangle inequality that the modulus of $\mathcal{G}_{\mathsf{TV}}$. is bounded. We illustrate this idea in Figure 4 and make it precise below.

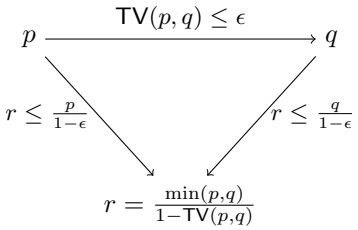

$$p \in \mathcal{G}_{\mathsf{TV}} \implies \|\mu(p) - \mu(r)\| \le \rho$$
$$q \in \mathcal{G}_{\mathsf{TV}} \implies \|\mu(q) - \mu(r)\| \le \rho$$
$$\implies \|\mu(p) - \mu(q)\| \le 2\rho$$

Figure 4: Midpoint distribution r helps bound the modulus for $\mathcal{G}_{\mathsf{TV}}$.

Recall that

$$\mathfrak{m}(\mathcal{G}_{\mathsf{TV}}(\rho,\epsilon), 2\epsilon) = \sup_{p,q \in \mathcal{G}_{\mathsf{TV}}(\rho,\epsilon) : \mathsf{TV}(p,q) \le 2\epsilon} \|\mu(p) - \mu(q)\|. \tag{10}$$

From $\mathsf{TV}(p,q) \le 2\epsilon$, we know that $r = \frac{\min(p,q)}{1-\mathsf{TV}(p,q)}$ can be obtained from either p and q by conditioning on an event of probability $1-\epsilon$. It then follows from $p, q \in \mathcal{G}_{\mathsf{TV}}(\rho,\epsilon)$ that $\|\mu(p) - \mu(r)\| \le \epsilon$ and similarly $\|\mu(q) - \mu(r)\| \le \epsilon$. Thus by the triangle inequality $\|\mu(p) - \mu(q)\| \le 2\rho$, which yields the desired result. □

We have seen so far that resilient distributions have bounded modulus, and that both Gaussian and bounded covariance distributions are resilient. The bound on the modulus for $\mathcal{G}_{\mathsf{cov}}$ that is implied by resilience is optimal ($\mathcal{O}(\sigma\sqrt{\epsilon})$), while for $\mathcal{G}_{\mathsf{gauss}}$ it is optimal up to log factors ($\mathcal{O}(\epsilon\sqrt{\log(1/\epsilon)})$ vs. $\mathcal{O}(\epsilon)$). In fact, Gaussians are a special case and resilience yields an essentially optimal bound at least for most non-parametric families of distributions. As one family of examples, consider distributions with bounded *Orlicz norm*:

Definition 2.11 (Orlicz norm)**.** A function $\psi : \mathbb{R}_{\ge 0} \to \mathbb{R}_{\ge 0}$ is an *Orlicz function* if ψ is convex, non-decreasing, and satisfies $\psi(0) = 0$, $\psi(x) \to \infty$ as $x \to \infty$. For an Orlicz function ψ, the Orlicz norm or ψ-norm of a random variable X is defined as

$$\|X\|_\psi \triangleq \inf\left\{t > 0 : \mathbb{E}_p\left[\psi\left(\frac{|X|}{t}\right)\right] \le 1\right\}. \tag{11}$$

We let $\mathcal{G}_\psi(\sigma)$ denote the family of distributions with $\|X - \mathbb{E}[X]\|_\psi \le \sigma$.

As special cases, we say that a random variable $X \sim p$ is *sub-Gaussian* with parameter σ if $\|\langle X - \mathbb{E}_p[X], v\rangle\|_{\psi_2} \le \sigma$ whenever $\|v\|_2 \le 1$, where $\psi_2(x) = e^{x^2} - 1$. We define a *sub-exponential* random variable similarly for the function $\psi_1(x) = e^x - 1$.

Definition 2.11 applies to distributions on $\mathbb{R}$, but we can generalize this to distributions on $\mathbb{R}^d$ by taking one-dimensional projections:

Definition 2.12 (Orlicz norm in $\mathbb{R}^d$)**.** For a random variable $X \in \mathbb{R}^d$ and Orlicz function ψ, we define the d-dimensional ψ-norm as

$$\|X\|_\psi \triangleq \inf\{t > 0 : \|\langle X, v\rangle\|_\psi \le t \text{ whenever } \|v\|_2 \le 1\}. \tag{12}$$

We let $\mathcal{G}_\psi(\sigma)$ denote the distributions with bounded ψ-norm as in Definition 2.11.

Thus a distribution has bounded ψ-norm if each of its 1-dimensional projections does. As an example, $\mathcal{G}_{\mathsf{cov}}(\sigma) = \mathcal{G}_\psi(\sigma)$ for $\psi(x) = x^2$, so Orlicz norms generalize bounded covariance. It is also possible to generalize Definition 2.12 to norms other than the ℓ_2-norm, which we will see in an exercise.

Functions with bounded Orlicz norm are resilient:

Lemma 2.13. *The family $\mathcal{G}_\psi(\sigma)$ is contained in $\mathcal{G}_{\mathsf{TV}}(2\sigma\epsilon\psi^{-1}(1/\epsilon), \epsilon)$ for all $0 < \epsilon < 1/2$.*

Proof. Without loss of generality assume $\mathbb{E}[X] = 0$. For any event E with $p(E) = 1 - \epsilon' \geq 1 - \epsilon$, denote its complement as E^c. We then have

$$\|\mathbb{E}_{X\sim p}[X \mid E]\|_2 \overset{(i)}{=} \frac{\epsilon'}{1-\epsilon'}\|\mathbb{E}_{X\sim p}[X \mid E^c]\|_2 \tag{13}$$

$$= \frac{\epsilon'}{1-\epsilon'} \sup_{\|v\|_2\leq 1} \mathbb{E}_{X\sim p}[\langle X, v\rangle \mid E^c] \tag{14}$$

$$\overset{(ii)}{\leq} \frac{\epsilon'}{1-\epsilon'} \sup_{\|v\|_2\leq 1} \sigma\psi^{-1}(\mathbb{E}_{X\sim p}[\psi(|\langle X, v\rangle|/\sigma) \mid E^c]) \tag{15}$$

$$\overset{(iii)}{\leq} \frac{\epsilon'}{1-\epsilon'} \sup_{\|v\|_2\leq 1} \sigma\psi^{-1}(\mathbb{E}_{X\sim p}[\psi(|\langle X, v\rangle|/\sigma)]/\epsilon') \tag{16}$$

$$\overset{(iv)}{\leq} \frac{\epsilon'}{1-\epsilon'}\sigma\psi^{-1}(1/\epsilon') \leq 2\epsilon\sigma\psi^{-1}(1/\epsilon), \tag{17}$$

as was to be shown. Here (i) is because $(1-\epsilon')\mathbb{E}[X \mid E] + \epsilon'\mathbb{E}[X \mid E^c] = 0$. Meanwhile (ii) is by convexity of ψ, (iii) is by non-negativity of ψ, and (iv) is the assumed ψ-norm bound. □

As a consequence, the modulus $\mathfrak{m}$ of $\mathcal{G}_\psi(\sigma)$ is $\mathcal{O}(\sigma\epsilon\psi^{-1}(1/\epsilon))$, and hence the minimum distance functional estimates the mean with error $\mathcal{O}(\sigma\epsilon\psi^{-1}(1/\epsilon))$. Note that for $\psi(x) = x^2$ this reproduces our result for bounded covariance. For $\psi(x) = x^k$ we get error $\mathcal{O}(\sigma\epsilon^{1-1/k})$ when a distribution has kth moments bounded by σ^k. Similarly for sub-Gaussian distributions we get error $\mathcal{O}(\sigma\epsilon\sqrt{\log(1/\epsilon)})$. We will show in an exercise that the error bound implied by Lemma 2.13 is optimal for any Orlicz function ψ.

Further properties and dual norm perspective. Having seen several examples of resilient distributions, we now collect some basic properties of resilience, as well as a dual perspective that is often fruitful. First, we can make the connection between resilience and tails even more precise with the following lemma:

Lemma 2.14. *For a fixed vector v, let $\tau_\epsilon(v)$ denote the ϵ-quantile of $\langle x - \mu, v\rangle$: $\mathbb{P}_{x\sim p}[\langle x - \mu, v\rangle \geq \tau_\epsilon(v)] = \epsilon$. Then, p is (ρ, ϵ)-resilient in a norm $\|\cdot\|$ if and only if the ϵ-tail of p has bounded mean when projected onto any dual unit vector v:*

$$\mathbb{E}_p[\langle x - \mu, v\rangle \mid \langle x - \mu, v\rangle \geq \tau_\epsilon(v)] \leq \frac{1-\epsilon}{\epsilon}\rho \text{ whenever } \|v\|_* \leq 1. \tag{18}$$

In particular, the ϵ-quantile satisfies $\tau_\epsilon(v) \leq \frac{1-\epsilon}{\epsilon}\rho$.

In other words, if we project onto any unit vector v in the dual norm, the ϵ-tail of $x - \mu$ must have mean at most $\frac{1-\epsilon}{\epsilon}\rho$. Lemma 2.14 is proved in Section C.

The intuition for Lemma 2.14 is the following picture, which is helpful to keep in mind more generally:

Specifically, letting $\hat{\mu} = \mathbb{E}[X \mid E]$, if we have $\|\hat{\mu} - \mu\| = \rho$, then there must be some dual norm unit vector v such that $\langle\hat{\mu} - \mu, v\rangle = \rho$ and $\|v\|_* = 1$. Moreover, for such a v, $\langle\hat{\mu} - \mu, v\rangle$ will be largest when E consists of the $(1-\epsilon)$-fraction of points for which $\langle X - \mu, v\rangle$ is largest. Therefore, resilience reduces to a 1-dimensional problem along each of the dual unit vectors v.

A related result establishes that for $\epsilon = \frac{1}{2}$, resilience in a norm is equivalent to having bounded first moments in the dual norm (see Section D for a proof):

Lemma 2.15. *Suppose that p is $(\rho, \frac{1}{2})$-resilient in a norm $\|\cdot\|$, and let $\|\cdot\|_*$ be the dual norm. Then p has 1st moments bounded by 2ρ: $\mathbb{E}_{x\sim p}[|\langle x - \mu, v\rangle|] \leq 2\rho\|v\|_*$ for all $v \in \mathbb{R}^d$.*

Conversely, if p has 1st moments bounded by ρ, it is $(2\rho, \frac{1}{2})$-resilient.

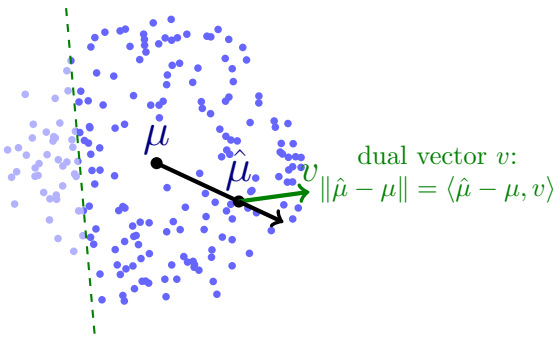

Figure 5: The optimal set T discards the smallest $\epsilon|S|$ elements projected onto a dual unit vector v.

Recap. We saw that the error of the trimmed mean grew as $\sqrt{d}$ in d dimensions, and introduced an alternative estimator–the minimum distance functional–that achieves better error. Specifically, it achieves error 2ρ for the family of (ρ, ϵ)-resilient distributions, which includes all distributions with bounded Orlicz norm (including bounded covariance, bounded moments, and sub-Gaussians).

The definition of resilience is important not just as an analysis tool. Without it, we would need a different estimator for each of the cases of bounded covariance, sub-Gaussian, etc., since the minimum distance functional depends on the family $\mathcal{G}$. Instead, we can always project onto the resilient family $\mathcal{G}_{\mathsf{TV}}(\rho, \epsilon)$ and be confident that this will typically yield an optimal error bound. The only complication is that projection still depends on the parameters ρ and ϵ; however, we can do without knowledge of either one of the parameters as long as we know the other.

[Lecture 3]

2.5 Concentration Inequalities

So far we have only considered the infinite-data limit where we directly observe $\tilde{p}$; but in general we would like to analyze what happens in finite samples where we only observe $X_1, \ldots, X_n$ sampled independently from $\tilde{p}$. In order to do this, we will want to be able to formalize statements such as "if we take the average of a large number of samples, it converges to the population mean". In order to do this, we will need a set of mathematical tools called *concentration inequalities*. A proper treatment of concentration could itself occupy an entire course, but we will cover the ideas here that are most relevant for our later analyses. See Boucheron et al. (2003), Boucheron et al. (2013), or Ledoux (2001) for more detailed expositions. Terence Tao also has some well-written lectures notes.

Concentration inequalities usually involve two steps:

1. We establish concentration for a single random variable, in terms of some property of that random variable.

2. We show that the property composes nicely for products of independent random variables.

A prototypical example (covered below) is showing that (1) a random variable has at most a $1/t^2$ probability of being t standard deviations from its mean; and (2) the standard deviation of a sum of n i.i.d. random variables is $\sqrt{n}$ times the standard deviation of a single variable.

The simplest concentration inequality is *Markov's inequality*. Consider the following question:

> A slot machine has an expected pay-out of \$5 (and its payout is always non-negative). What can we say about the probability that it pays out at least \$100?

We observe that the probability must be at most 0.05, since a 0.05 chance of a \$100 payout would by itself already contribute \$5 to the expected value. Moreover, this bound is achievable by taking a slot machine that pays \$0 with probability 0.95 and \$100 with probability 0.05. Markov's inequality is the generalization of this observation:

Theorem 2.16 (Markov's inequality)**.** *Let X be a non-negative random variable with mean μ. Then, $\mathbb{P}[X \geq t \cdot \mu] \leq \frac{1}{t}$.*

Markov's inequality accomplishes our first goal of establishing concentration for a single random variable, but it has two issues: first, the $\frac{1}{t}$ tail bound decays too slowly in many cases (we instead would like exponentially decaying tails); second, Markov's inequality doesn't compose well and so doesn't accomplish our second goal.

We can address both issues by applying Markov's inequality to some transformed random variable. For instance, applying Markov's inequality to the random variable $Z = (X - \mu)^2$ yields the stronger *Chebyshev inequality*:

Theorem 2.17 (Chebyshev's inequality)**.** *Let X be a real-valued random variable with mean μ and variance σ^2. Then, $\mathbb{P}[|X - \mu| \geq t \cdot \sigma] \leq \frac{1}{t^2}$.*

Proof. Since $Z = (X - \mu)^2$ is non-negative, we have that $\mathbb{P}[Z \geq t^2 \cdot \sigma^2] \leq \frac{1}{t^2}$ by Markov's inequality. Taking the square-root gives $\mathbb{P}[|X - \mu| \geq t \cdot \sigma] \leq \frac{1}{t^2}$, as was to be shown. □

Chebyshev's inequality improves the $1/t$ dependence to $1/t^2$. But more importantly, it gives a bound in terms of a quantity (the variance σ^2) that composes nicely:

Lemma 2.18 (Additivity of variance)**.** *Let $X_1, \ldots, X_n$ be pairwise independent random variables, and let $\mathsf{Var}[X]$ denote the variance of X. Then,*

$$\mathsf{Var}[X_1 + \cdots + X_n] = \mathsf{Var}[X_1] + \cdots + \mathsf{Var}[X_n]. \tag{19}$$

Proof. It suffices by induction to prove this for two random variables. Without loss of generality assume that both variables have mean zero. Then we have $\mathsf{Var}[X + Y] = \mathbb{E}[(X + Y)^2] = \mathbb{E}[X^2] + \mathbb{E}[Y^2] + 2\mathbb{E}[XY] = \mathsf{Var}[X] + \mathsf{Var}[Y] + 2\mathbb{E}[X]\mathbb{E}[Y] = \mathsf{Var}[X] + \mathsf{Var}[Y]$, where the second-to-last step uses pairwise independence. □

Chebyshev's inequality together with Lemma 2.18 together allow us to show that an average of i.i.d. random variables converges to its mean at a $1/\sqrt{n}$ rate:

Corollary 2.19. *Suppose $X_1, \ldots, X_n$ are drawn i.i.d. from p, where p has mean μ and variance σ^2. Also let $S = \frac{1}{n}(X_1 + \cdots + X_n)$. Then, $\mathbb{P}[|S - \mu|/\sigma \geq t/\sqrt{n}] \leq 1/t^2$.*

Proof. Lemma 2.18 implies that $\mathsf{Var}[S] = \sigma^2/n$, from which the result follows by Chebyshev's inequality. □

Higher moments. Chebyshev's inequality gives bounds in terms of the second moment of $X - \mu$. Can we do better by considering higher moments such as the 4th moment? Supposing that $\mathbb{E}[(X - \mu)^4] \leq \tau^4$, we do get the analogous bound $\mathbb{P}[|X - \mu| \geq t \cdot \tau] \leq 1/t^4$. However, the 4th moment doesn't compose as nicely as the variance; if X and Y are two independent mean-zero random variables, then we have

$$\mathbb{E}[(X + Y)^4] = \mathbb{E}[X^4] + \mathbb{E}[Y^4] + 6\mathbb{E}[X^2]\mathbb{E}[Y^2], \tag{20}$$

where the $\mathbb{E}[X^2]\mathbb{E}[Y^2]$ can't be easily dealt with. It is possible to bound higher moments under composition, for instance using the *Rosenthal inequality* which states that

$$\mathbb{E}[|\sum_i X_i|^p] \leq \mathcal{O}(p)^p \sum_i \mathbb{E}[|X_i|^p] + \mathcal{O}(\sqrt{p})^p (\sum_i \mathbb{E}[X_i^2])^{p/2} \tag{21}$$

for independent random variables X_i. Note that the first term on the right-hand-side typically grows as $n \cdot \mathcal{O}(p)^p$ while the second term typically grows as $\mathcal{O}(\sqrt{pn})^p$.

We will typically not take the Rosenthal approach and instead work with an alternative, nicer object called the *moment generating function*:

$$m_X(\lambda) \stackrel{\text{def}}{=} \mathbb{E}[\exp(\lambda(X - \mu))]. \tag{22}$$

For independent random variables, the moment generating function composes via the identity $m_{X_1 + \cdots + X_n}(\lambda) = \prod_{i=1}^n m_{X_i}(\lambda)$. Applying Markov's inequality to the moment generating function yields the *Chernoff bound*:

Theorem 2.20 (Chernoff bound)**.** *For a random variable X with moment generating $m_X(\lambda)$, we have*

$$\mathbb{P}[X - \mu \geq t] \leq \inf_{\lambda \geq 0} m_X(\lambda) e^{-\lambda t}. \tag{23}$$

Proof. By Markov's inequality, $\mathbb{P}[X - \mu \geq t] = \mathbb{P}[\exp(\lambda(X - \mu)) \geq \exp(\lambda t)] \leq \mathbb{E}[\exp(\lambda(X - \mu))] / \exp(\lambda t)$, which is equal to $m_X(\lambda) e^{-\lambda t}$ by the definition of m_X. Taking inf over λ yields the claimed bound. □

Sub-exponential and sub-Gaussian distributions. An important special case is sub-exponential random variables; recall these are random variables satisfying $\mathbb{E}[\exp(|X - \mu|/\sigma)] \leq 2$. For these, applying the Chernoff bound with $\lambda = 1/\sigma$ yields $\mathbb{P}[X - \mu \geq t] \leq 2e^{-t/\sigma}$.

Another special case is sub-Gaussian random variables (those satisfying $\mathbb{E}[\exp((X - \mu)^2/\sigma^2)] \leq 2$). In this case, using the inequality $ab \leq a^2/4 + b^2$, we have

$$m_X(\lambda) = \mathbb{E}[\exp(\lambda(X - \mu))] \leq \mathbb{E}[\exp(\lambda^2\sigma^2/4 + (X - \mu)^2/\sigma^2)] \leq 2\exp(\lambda^2\sigma^2/4). \tag{24}$$

The factor of 2 is pesky and actually we can get the more convenient bound $m_X(\lambda) \leq \exp(3\lambda^2\sigma^2/2)$ (Rivasplata, 2012). Plugging this into the Chernoff bound yields $\mathbb{P}[X - \mu \geq t] \leq \exp(3\lambda^2\sigma^2/2 - \lambda t)$; minimizing over λ gives the optimized bound $\mathbb{P}[X - \mu \geq t] \leq \exp(-t^2/6\sigma^2)$.

Sub-Gaussians are particularly convenient because the bound $m_X(\lambda) \leq \exp(3\lambda^2\sigma^2/2)$ composes well. Let $X_1, \ldots, X_n$ be independent sub-Gaussians with constants $\sigma_1, \ldots, \sigma_n$. Then we have $m_{X_1 + \cdots + X_n}(\lambda) \leq \exp(3\lambda^2(\sigma_1^2 + \cdots + \sigma_n^2)/2)$. We will use this to bound the behavior of sums of bounded random variables using *Hoeffding's inequality*:[1]

Theorem 2.21 (Hoeffding's inequality)**.** *Let $X_1, \ldots, X_n$ be zero-mean random variables lying in $[-M, M]$, and let $S = \frac{1}{n}(X_1 + \cdots + X_n)$. Then, $\mathbb{P}[S \geq t] \leq \exp(-\ln(2)nt^2/6M^2) \leq \exp(-nt^2/9M^2)$.*

Proof. First, note that each X_i is sub-Gaussian with parameter $\sigma = M/\sqrt{\ln 2}$, since $\mathbb{E}[\exp(X_i^2/\sigma^2)] \leq \exp(M^2/\sigma^2) = \exp(\ln(2)) = 2$. We thus have $m_{X_i}(\lambda) \leq \exp(3\lambda^2 M^2/2\ln 2)$, and so by the multiplicativity of moment generating functions we obtain $m_S(\lambda) \leq \exp(3\lambda^2 M^2/(2n\ln 2))$. Plugging into Chernoff's bound and optimizing λ as before yields $\mathbb{P}[S \geq t] \leq \exp(-\ln(2)nt^2/6M^2)$ as claimed. □

Hoeffding's inequality shows that a sum of independent random variables converges to its mean at a $1/\sqrt{n}$ rate, with tails that decay as fast as a Gaussian as long as each of the individual variables is bounded. Compare this to the $1/t^2$ decay that we obtained earlier through Chebyshev's inequality.

Cumulants. The moment generating function is a convenient tool because it multiplies over independent random variables. However, its existence requires that X already have thin tails, since $\mathbb{E}[\exp(\lambda X)]$ must be finite. For heavy-tailed distributions a (laborious) alternative is to use *cumulants*.

The cumulant function is defined as

$$K_X(\lambda) \stackrel{\text{def}}{=} \log \mathbb{E}[\exp(\lambda X)]. \tag{25}$$

Note this is the log of the moment-generating function. Taking the log is convenient because now we have additivity: $K_{X+Y}(\lambda) = K_X(\lambda) + K_Y(\lambda)$ for independent X, Y. Cumulants are obtained by writing $K_X(\lambda)$ as a power series:

$$K_X(\lambda) = 1 + \sum_{n=1}^{\infty} \frac{\kappa_n(X)}{n!} \lambda^n. \tag{26}$$

[1] Most of the constants presented here are suboptimal; we have focused on giving simpler proofs at the expense of sharp constants.

When $\mathbb{E}[X] = 0$, the first few values of κ_n are:

$$\kappa_1(X) = 0, \tag{27}$$
$$\kappa_2(X) = \mathbb{E}[X^2], \tag{28}$$
$$\kappa_3(X) = \mathbb{E}[X^3], \tag{29}$$
$$\kappa_4(X) = \mathbb{E}[X^4] - 3\mathbb{E}[X^2]^2, \tag{30}$$
$$\kappa_5(X) = \mathbb{E}[X^5] - 10\mathbb{E}[X^3]\mathbb{E}[X^2], \tag{31}$$
$$\kappa_6(X) = \mathbb{E}[X^6] - 16\mathbb{E}[X^4]\mathbb{E}[X^2] - 10\mathbb{E}[X^3]^2 + 30\mathbb{E}[X^2]^3. \tag{32}$$

Since K is additive, each of the κ_n are as well. Thus while we ran into the issue that $\mathbb{E}[(X+Y)^4] \neq \mathbb{E}[X^4] + \mathbb{E}[Y^4]$, it is the case that $\kappa_4(X+Y) = \kappa_4(X) + \kappa_4(Y)$ as long as X and Y are independent. By going back and forth between moments and cumulants it is possible to obtain tail bounds even if only some of the moments exist. However, this can be arduous and Rosenthal's inequality is probably the better route in such cases.

[Lecture 4]

2.5.1 Applications of concentration inequalities

Having developed the machinery above, we next apply it to a few concrete problems to give a sense of how to use it. A key lemma which we will use repeatedly is the union bound, which states that if $E_1, \dots, E_n$ are events with probabilities $\pi_1, \dots, \pi_n$, then the probability of $E_1 \cup \cdots \cup E_n$ is at most $\pi_1 + \cdots + \pi_n$. A corollary is that if n events each have probability $\ll 1/n$, then there is a large probability that none of the events occur.

Maximum of sub-Gaussians. Suppose that $X_1, \dots, X_n$ are mean-zero sub-Gaussian with parameter σ, and let $Y = \max_{i=1}^n X_i$. How large is Y? We will show the following:

Lemma 2.22. *The random variable Y is $\mathcal{O}(\sigma\sqrt{\log(n/\delta)})$ with probability $1-\delta$.*

Proof. By the Chernoff bound for sub-Gaussians, we have that $\mathbb{P}[X_i \geq \sigma\sqrt{6\log(n/\delta)}] \leq \exp(-\log(n/\delta)) = \delta/n$. Thus by the union bound, the probability that any of the X_i exceed $\sigma\sqrt{6\log(n/\delta)}$ is at most δ. Thus with probability at least $1-\delta$ we have $Y \leq \sigma\sqrt{6\log(n/\delta)}$, as claimed. □

Lemma 2.22 illustrates a typical proof strategy: We first decompose the event we care about as a union of simpler events, then show that each individual event holds with high probability by exploiting independence. As long as the "failure probability" of a single event is much small than the inverse of the number of events, we obtain a meaningful bound. In fact, this strategy can be employed even for an infinite number of events by discretizing to an "ϵ-net", as we will see below:

Eigenvalue of random matrix. Let $X_1, \dots, X_n$ be independent zero-mean sub-Gaussian variables in $\mathbb{R}^d$ with parameter σ, and let $M = \frac{1}{n}\sum_{i=1}^n X_i X_i^\top$. How large is $\|M\|$, the maximum eigenvalue of M? We will show:

Lemma 2.23. *The maximum eigenvalue $\|M\|$ is $\mathcal{O}(\sigma^2 \cdot (1 + d/n + \log(1/\delta)/n))$ with probability $1-\delta$.*

Proof. The maximum eigenvalue can be expressed as

$$\|M\| = \sup_{\|v\|_2 \leq 1} v^\top M v = \sup_{\|v\|_2 \leq 1} \frac{1}{n}\sum_{i=1}^n |\langle X_i, v\rangle|^2. \tag{33}$$

The quantity inside the sup is attractive to analyze because it is an average of independent random variables. Indeed, we have

$$\mathbb{E}[\exp(\frac{n}{\sigma^2}v^\top Mv)] = \mathbb{E}[\exp(\sum_{i=1}^n |\langle X_i, v\rangle|^2/\sigma^2)] \tag{34}$$

$$= \prod_{i=1}^n \mathbb{E}[\exp(|\langle X_i, v\rangle|^2/\sigma^2)] \le 2^n, \tag{35}$$

where the last step follows by sub-Gaussianity if $\langle X_i, v\rangle$. The Chernoff bound then gives $\mathbb{P}[v^\top Mv \ge t] \le 2^n \exp(-nt/\sigma^2)$.

If we were to follow the same strategy as Lemma 2.22, the next step would be to union bound over v. Unfortunately, there are infinitely many v so we cannot do this directly. Fortunately, we can get by with only considering a large but finite number of v; we will construct a finite subset $\mathcal{N}_{1/4}$ of the unit ball such that

$$\sup_{v\in\mathcal{N}_{1/4}} v^\top Mv \ge \frac{1}{2}\sup_{\|v\|_2\le 1} v^\top Mv. \tag{36}$$

Our construction follows Section 5.2.2 of Vershynin (2010). Let $\mathcal{N}_{1/4}$ be a maximal set of points in the unit ball such that $\|x-y\|_2 \ge 1/4$ for all distinct $x, y \in \mathcal{N}_{1/4}$. We observe that $|\mathcal{N}_{1/4}| \le 9^d$; this is because the balls of radius $1/8$ around each point in $\mathcal{N}_{1/4}$ are disjoint and contained in a ball of radius $9/8$.

To establish (36), let v maximize $v^\top Mv$ over $\|v\|_2 \le 1$ and let u maximize $v^\top Mv$ over $\mathcal{N}_{1/4}$. Then

$$|v^\top Mv - u^\top Mu| = |v^\top M(v-u) + u^\top M(v-u)| \tag{37}$$

$$\le (\|v\|_2 + \|u\|_2)\|M\|\|v-u\|_2 \tag{38}$$

$$\le 2\cdot\|M\|\cdot(1/4) = \|M\|/2. \tag{39}$$

Since $v^\top Mv = \|M\|$, we obtain $|\|M\| - u^\top Mu| \le \|M\|/2$, whence $u^\top Mu \ge \|M\|/2$, which establishes (36). We are now ready to apply the union bound: Recall that from the Chernoff bound on $v^\top Mv$, we had $\mathbb{P}[v^\top Mv \ge t] \le 2^n\exp(-nt/\sigma^2)$, so

$$\mathbb{P}[\sup_{v\in\mathcal{N}_{1/4}} v^\top Mv \ge t] \le 9^d 2^n \exp(-nt/\sigma^2). \tag{40}$$

Solving for this quantity to equal δ, we obtain

$$t = \frac{\sigma^2}{n}\cdot(n\log(2) + d\log(9) + \log(1/\delta)) = \mathcal{O}(\sigma^2\cdot(1 + d/n + \log(1/\delta)/n)), \tag{41}$$

as was to be shown. □

VC dimension. Our final example will be important in the following section; it concerns how quickly a family of events with certain geometric structure converges to its expectation. Let $\mathcal{H}$ be a collection of functions $f : \mathcal{X} \to \{0,1\}$, and define the *VC dimension* $\mathsf{vc}(\mathcal{H})$ to be the maximum d for which there are points $x_1, \ldots, x_d$ such that $(f(x_1), \ldots, f(x_d))$ can take on all 2^d possible values. For instance:

- If $\mathcal{X} = \mathbb{R}$ and $\mathcal{H} = \{\mathbb{I}[x \ge \tau] \mid \tau \in \mathbb{R}\}$ is the family of threshold functions, then $\mathsf{vc}(\mathcal{H}) = 1$.
- If $\mathcal{X} = \mathbb{R}^d$ and $\mathcal{H} = \{\mathbb{I}[\langle x, v\rangle \ge \tau] \mid v \in \mathbb{R}^d, \tau \in \mathbb{R}\}$ is the family of half-spaces, then $\mathsf{vc}(\mathcal{H}) = d+1$.

Additionally, for a point set $S = \{x_1, \ldots, x_n\}$, let $V_\mathcal{H}(S)$ denote the number of distinct values of $(f(x_1), \ldots, f(x_n))$ and $V_\mathcal{H}(n) = \max\{V_\mathcal{H}(S) \mid |S| = n\}$. Thus the VC dimension is exactly the maximum n such that $V_\mathcal{H}(n) = 2^n$.

We will show the following:

Proposition 2.24. *Let $\mathcal{H}$ be a family of functions with $\mathsf{vc}(H) = d$, and let $X_1, \ldots, X_n \sim p$ be i.i.d. random variables over $\mathcal{X}$. For $f : \mathcal{X} \to \{0,1\}$, let $\nu_n(f) = \frac{1}{n}|\{i \mid f(X_i) = 1\}|$ and let $\nu(f) = p(f(X) = 1)$. Then*

$$\sup_{f\in\mathcal{H}} |\nu_n(f) - \nu(f)| \le \mathcal{O}\Big(\sqrt{\frac{d + \log(1/\delta)}{n}}\Big) \tag{42}$$

with probability $1-\delta$.

We will prove a weaker result that has a $d\log(n)$ factor instead of d, and which bounds the expected value rather than giving a probability $1-\delta$ bound. The $\log(1/\delta)$ tail bound follows from *McDiarmid's inequality*, which is a standard result in a probability course but requires tools that would take us too far afield. Removing the $\log(n)$ factor is slightly more involved and uses a tool called *chaining*.

Proof of Proposition 2.24. The importance of the VC dimension for our purposes lies in the Sauer-Shelah lemma:

Lemma 2.25 (Sauer-Shelah)**.** *Let $d = \mathsf{vc}(\mathcal{H})$. Then $V_{\mathcal{H}}(n) \le \sum_{k=0}^{d} \binom{n}{k} \le 2n^d$.*

It is tempting to union bound over the at most $V_{\mathcal{H}}(n)$ distinct values of $(f(X_1),\dots,f(X_n))$; however, this doesn't work because revealing $X_1,\dots,X_n$ uses up all of the randomness in the problem and we have no randomness left from which to get a concentration inequality! We will instead have to introduce some new randomness using a technique called *symmetrization*.

Regarding the expectation, let $X_1',\dots,X_n'$ be independent copies of $X_1,\dots,X_n$ and let $\nu_n'(f)$ denote the version of $\nu_n(f)$ computed with the X_i'. Then we have

$$\mathbb{E}_X[\sup_{f\in\mathcal{H}} |\nu_n(f)-\nu(f)|] \le \mathbb{E}_{X,X'}[\sup_{f\in\mathcal{H}} |\nu_n(f)-\nu_n'(f)|] \tag{43}$$

$$= \frac{1}{n}\mathbb{E}_{X,X'}[\sup_{f\in\mathcal{H}} |\sum_{i=1}^{n} f(X_i)-f(X_i')|. \tag{44}$$

We can create our new randomness by noting that since X_i and X_i' are identically distributed, $f(X_i)-f(X_i')$ has the same distribution as $s_i(f(X_i)-f(X_i'))$, where s_i is a random sign variable that is ±1 with equal probability. Introducing these variables and continuing the inequality, we thus have

$$\frac{1}{n}\mathbb{E}_{X,X'}[\sup_{f\in\mathcal{H}} |\sum_{i=1}^{n} f(X_i)-f(X_i')|] = \frac{1}{n}\mathbb{E}_{X,X',s}[\sup_{f\in\mathcal{H}} |\sum_{i=1}^{n} s_i(f(X_i)-f(X_i'))|]. \tag{45}$$

We now have enough randomness to exploit the Sauer-Shelah lemma. If we fix X and X', note that the quantities $f(X_i)-f(X_i')$ take values in $[-1,1]$ and collectively can take on at most $V_{\mathcal{H}}(n)^2 = \mathcal{O}(n^{2d})$ values. But for fixed X, X', the quantities $s_i(f(X_i)-f(X_i'))$ are independent, zero-mean, bounded random variables and hence for fixed f we have $\mathbb{P}[\sum_i s_i(f(X_i)-f(X_i')) \ge t] \le \exp(-t^2/9n)$ by Hoeffding's inequality. Union bounding over the $\mathcal{O}(n^{2d})$ effectively distinct f, we obtain

$$\mathbb{P}_s[\sup_{f\in\mathcal{H}} |\sum_{i} s_i(f(X_i)-f(X_i'))| \ge t \mid X, X'] \le \mathcal{O}(n^{2d})\exp(-t^2/9n). \tag{46}$$

This is small as long as $t \gg \sqrt{nd\log n}$, so (45) is $\mathcal{O}(\sqrt{d\log n/n})$, as claimed. □

A particular consequence of Proposition 2.24 is the *Dvoretzky-Kiefer-Wolfowitz inequality*:

Proposition 2.26 (DKW inequality)**.** *For a distribution p on $\mathbb{R}$ and i.d.d. samples $X_1,\dots,X_n \sim p$, define the empirical cumulative density function as $F_n(x) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{I}[X_i \le x]$, and the population cumulative density function as $F(x) = p(X \le x)$. Then $\mathbb{P}[\sup_{x\in\mathbb{R}} |F_n(x)-F(x)| \ge t] \le 2e^{-2nt^2}$.*

This follows from applying Proposition 2.24 to the family of threshold functions.

[Lecture 5]

2.6 Finite-Sample Analysis

Now that we have developed tools for analyzing statistical concentration, we will use these to analyze the finite-sample behavior of robust estimators. Recall that we previously studied the minimum distance functional defined as

$$\hat{\theta}(\tilde{p}) = \theta^*(q), \text{ where } q = \arg\min_{q\in\mathcal{G}} \mathsf{TV}(q,\tilde{p}). \tag{47}$$

This projects onto the set $\mathcal{G}$ under TV distance and outputs the optimal parameters for the projected distribution.

The problem with the minimum distance functional defined above is that projection under TV usually doesn't make sense for finite samples! For instance, suppose that p is a Gaussian distribution and let p_n and p'_n be the empirical distributions of two different sets of n samples. Then $\mathsf{TV}(p_n, p) = \mathsf{TV}(p_n, p'_n) = 1$ almost surely. This is because samples from a continuous probability distribution will almost surely be distinct, and TV distance doesn't give credit for being "close"—the TV distance between two point masses at 1 and 1.000001 is still 1.[2]

To address this issue, we will consider two solutions. The first solution is to *relax the distance*. Intuitively, the issue is that the TV distance is too strong—it reports a large distance even between a population distribution p and the finite-sample distribution p_n. We will replace the distance TV with a more forgiving distance $\widetilde{\mathsf{TV}}$ and use the minimum distance functional corresponding to this relaxed distance. To show that projection under $\widetilde{\mathsf{TV}}$ still works, we will need to check that the modulus $\mathfrak{m}(\mathcal{G}, \epsilon)$ is still small after we replace TV with $\widetilde{\mathsf{TV}}$, and we will also need to check that the distance $\widetilde{\mathsf{TV}}(p, p_n)$ between p and its empirical distribution is small with high probability. We do this below in Section 2.6.1.

An alternative to relaxing the distance from TV to $\widetilde{\mathsf{TV}}$ is to expand the destination set from $\mathcal{G}$ to some $\mathcal{M} \supset \mathcal{G}$, such that even though p is not close to the empirical distribution p_n, *some* element of $\mathcal{M}$ is close to p_n. Another advantage to expanding the destination set is that projecting onto $\mathcal{G}$ may not be computationally tractable, while projecting onto some larger set $\mathcal{M}$ can sometimes be done efficiently. We will see how to statistically analyze this modified projection algorithm in Section 2.6.2, and study the computational feasibility of projecting onto a set $\mathcal{M}$ starting in Section 2.7.

2.6.1 Relaxing the Distance

Here we instantiate the first solution of replacing TV with some $\widetilde{\mathsf{TV}}$ for the projection algorithm. The following lemma shows that properties we need $\widetilde{\mathsf{TV}}$ to satisfy:

Lemma 2.27. *Suppose that $\widetilde{\mathsf{TV}}$ is a (pseudo-)metric such that $\widetilde{\mathsf{TV}}(p, q) \le \mathsf{TV}(p, q)$ for all p, q. If we assume that $p^* \in \mathcal{G}$ and $\mathsf{TV}(p^*, \tilde{p}) \le \epsilon$, then the error of the minimum distance functional (2) with $D = \widetilde{\mathsf{TV}}$ is at most $\mathfrak{m}(\mathcal{G}, 2\epsilon', \widetilde{\mathsf{TV}}, L)$, where $\epsilon' = \epsilon + \widetilde{\mathsf{TV}}(\tilde{p}, \tilde{p}_n)$.*

Proof. By Proposition 2.4 we already know that the error is bounded by $\mathfrak{m}(\mathcal{G}, 2\widetilde{\mathsf{TV}}(p^*, \tilde{p}_n), \widetilde{\mathsf{TV}}, L)$. Since $\widetilde{\mathsf{TV}}$ is a pseudometric, by the triangle inequality we have $\widetilde{\mathsf{TV}}(p^*, \tilde{p}_n) \le \widetilde{\mathsf{TV}}(p^*, \tilde{p}) + \widetilde{\mathsf{TV}}(\tilde{p}, \tilde{p}_n)$. Finally, $\widetilde{\mathsf{TV}}(p^*, \tilde{p}) \le \mathsf{TV}(p^*, \tilde{p})$ by assumption. □

Lemma 2.27 shows that we need $\widetilde{\mathsf{TV}}$ to satisfy two properties: $\widetilde{\mathsf{TV}}(\tilde{p}, \tilde{p}_n)$ should be small, and the modulus $\mathfrak{m}(\mathcal{G}, \epsilon, \widetilde{\mathsf{TV}})$ should not be too much larger than $\mathfrak{m}(\mathcal{G}, \epsilon, \mathsf{TV})$.

For mean estimation (where recall $L(p, \theta) = \|\theta - \mu(p)\|_2$), we will use the following $\widetilde{\mathsf{TV}}$:

$$\widetilde{\mathsf{TV}}_{\mathcal{H}}(p, q) \stackrel{\text{def}}{=} \sup_{f \in \mathcal{H}, \tau \in \mathbb{R}} |\mathbb{P}_{X \sim p}[f(X) \ge \tau] - \mathbb{P}_{X \sim q}[f(X) \ge \tau]|. \tag{48}$$

(Note the similarity to the distance in Proposition 2.24; we will make use of this later.) We will make the particular choice $\mathcal{H} = \mathcal{H}_{\mathsf{lin}}$, where $\mathcal{H}_{\mathsf{lin}} \stackrel{\text{def}}{=} \{x \mapsto \langle v, x \rangle \mid v \in \mathbb{R}^d\}$.

First note that $\widetilde{\mathsf{TV}}_{\mathcal{H}}$ is indeed upper-bounded by TV, since $\mathsf{TV}(p, q) = \sup_E |p(E) - q(E)|$ is the supremum over all events E, and (48) takes a supremum over a subset of events. The intuition for taking the particular family $\mathcal{H}$ is that linear projections of our data contain all information needed to recover the mean, so perhaps it is enough for distributions to be close only under these projections.

Bounding the modulus. To formalize this intuition, we prove the following *mean crossing lemma*:

Lemma 2.28. *Suppose that p and q are two distributions such that $\widetilde{\mathsf{TV}}_{\mathcal{H}}(p, q) \le \epsilon$. Then for any $f \in \mathcal{H}$, there are distributions $r_p \le \frac{p}{1-\epsilon}$ and $r_q \le \frac{q}{1-\epsilon}$ such that $\mathbb{E}_{X \sim r_p}[f(X)] \ge \mathbb{E}_{Y \sim r_q}[f(Y)]$.*

[2] We will later study the W_1 distance, which does give credit for being close.

Proof. We will prove the stronger statement that $f(X)$ under r_p *stochastically dominates* $f(Y)$ under r_q. Starting from p, q, we delete ϵ probability mass corresponding to the smallest points of $f(X)$ in p to get r_p, and delete ϵ probability mass corresponding to the largest points $f(Y)$ in q to get r_q. Since $\widetilde{\mathsf{TV}}_{\mathcal{H}}(p, q) \le \epsilon$ we have

$$\sup_{\tau \in \mathbb{R}} |\mathbb{P}_{X\sim p}(f(X) \ge \tau) - \mathbb{P}_{Y\sim q}(f(Y) \ge \tau)| \le \epsilon, \tag{49}$$

which implies that $\mathbb{P}_{r_p}(f(X) \ge \tau) \ge \mathbb{P}_{r_q}(f(Y) \ge \tau)$ for all $t \in \mathbb{R}$. Hence, r_p stochastically dominates r_q and $\mathbb{E}_{r_p}[f(X)] \ge \mathbb{E}_{r_q}[(Y)]$. □

Mean crossing lemmas such as Lemma 2.28 help us bound the modulus of relaxed distances for the family of resilient distributions. In this case we have the following corollary:

Corollary 2.29. *For the family $\mathcal{G}_{\mathsf{TV}}(\rho, \epsilon)$ of (ρ, ϵ)-resilient distributions and $L(p, \theta) = \|\theta - \mu(p)\|_2$, we have*

$$\mathfrak{m}(\mathcal{G}_{\mathsf{TV}}(\rho, \epsilon), \epsilon, \widetilde{\mathsf{TV}}_{\mathcal{H}_{\mathsf{lin}}}) \le 2\rho. \tag{50}$$

Compare to Theorem 2.10 where we showed that $\mathfrak{m}(\mathcal{G}_{\mathsf{TV}}, \epsilon, \mathsf{TV}) \le \rho$. Thus as long as Theorem 2.10 is tight, relaxing from TV to $\widetilde{\mathsf{TV}}_{\mathcal{H}_{\mathsf{lin}}}$ doesn't increase the modulus at all!

Proof of Corollary 2.29. Let $p, q \in \mathcal{G}_{\mathsf{TV}}$ such that $\widetilde{\mathsf{TV}}(p, q) \le \epsilon$. Take $v = \arg\max_{\|v\|_2=1} v^\top(\mathbb{E}_p[X] - \mathbb{E}_q[X])$, hence $\mathbb{E}_p[v^\top X] - \mathbb{E}_q[v^\top X] = \|\mathbb{E}_p[X] - \mathbb{E}_q[X]\|_2$. It follows from Lemma 2.28 that there exist $r_p \le \frac{p}{1-\epsilon}, r_q \le \frac{q}{1-\epsilon}$ such that

$$\mathbb{E}_{r_p}[v^\top X] \le \mathbb{E}_{r_q}[v^\top X]. \tag{51}$$

Furthermore, from $p, q \in \mathcal{G}_{\mathsf{TV}}(\rho, \epsilon)$, we have

$$\mathbb{E}_p[v^\top X] - \mathbb{E}_{r_p}[v^\top X] \le \rho, \tag{52}$$
$$\mathbb{E}_{r_q}[v^\top X] - \mathbb{E}_q[v^\top X] \le \rho. \tag{53}$$

Then,

$$\|\mathbb{E}_p[X] - \mathbb{E}_q[X]\|_2 = \mathbb{E}_p[v^\top X] - \mathbb{E}_q[v^\top X] \tag{54}$$
$$= \underbrace{\mathbb{E}_p[v^\top X] - \mathbb{E}_{r_p}[v^\top X]}_{\le \rho} + \underbrace{\mathbb{E}_{r_p}[v^\top X] - \mathbb{E}_{r_q}[v^\top X]}_{\le 0} + \underbrace{\mathbb{E}_{r_q}[v^\top X] - \mathbb{E}_q[v^\top X]}_{\le \rho} \tag{55}$$
$$\le 2\rho, \tag{56}$$

which shows the modulus is small as claimed. □

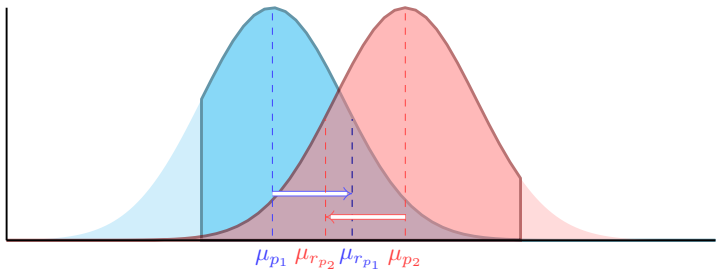

Figure 6: Illustration of mean cross lemma. For any distributions p_1, p_2 that are close under $\widetilde{\mathsf{TV}}$, we can truncate the ϵ-tails of each distribution to make their means cross.

Bounding the distance to the empirical distribution. Now that we have bounded the modulus, it remains to bound the distance $\widetilde{\mathsf{TV}}(\tilde{p}, \tilde{p}_n)$. Note that $\widetilde{\mathsf{TV}}(\tilde{p}, \tilde{p}_n)$ is exactly the quantity bounded in equation (42) of Proposition 2.24; we thus have that $\widetilde{\mathsf{TV}}_{\mathcal{H}}(\tilde{p}, \tilde{p}_n) \le \mathcal{O}\big(\sqrt{\frac{\mathsf{vc}(\mathcal{H})+\log(1/\delta)}{n}}\big)$ with probability $1-\delta$. Here $\mathsf{vc}(\mathcal{H})$ is the VC dimension of the family of threshold functions $\{x \mapsto \mathbb{I}[f(x) \ge \tau] \mid f \in \mathcal{H}, \tau \in \mathbb{R}\}$. So, for $\mathcal{H} = \mathcal{H}_{\mathsf{lin}}$ all we need to do is bound the VC dimension of the family of halfspace functions on $\mathbb{R}^d$.

We claimed earlier that this VC dimension is $d+1$, but we prove it here for completeness. We will show that no set of points $x_1, \dots, x_{d+2} \in \mathbb{R}^d$ cannot be shattered into all 2^{d+2} possible subsets using halfspaces. For any such points we can find multipliers $a_1, \dots, a_{d+2} \in \mathbb{R}$ such that

$$\sum_{i=1}^{d+2} a_i x_i = 0, \sum_{i=1}^{d+2} a_i = 0. \tag{57}$$

Let $S_+ = \{i \mid a_i > 0\}$ and $S_- = \{i \mid a_i < 0\}$. We will show that the convex hulls of S_+ and S_- intersect. Consequently, there is no vector v and threshold τ such that $\langle x_i, v\rangle \ge \tau$ iff $i \in S_+$. (This is because both a halfspace and its complement are convex, so if we let $H_{v,\tau}$ denote the half-space, it is impossible to have $S_+ \subset H_{v,\tau}$, $S_- \subset H_{v,\tau}^c$, and $\mathrm{conv}(S_+) \cap \mathrm{conv}(S_-) \ne \emptyset$.)

To prove that the convex hulls intersect, note that we have

$$\frac{1}{A}\sum_{i \in S_+} a_i x_i = \frac{1}{A}\sum_{i \in S_-} (-a_i) x_i, \tag{58}$$

where $A = \sum_{i \in S_+} a_i = \sum_{i \in S_-} (-a_i)$. But the left-hand-side lies in $\mathrm{conv}(S_+)$ while the right-hand-side lies in $\mathrm{conv}(S_-)$, so the convex hulls do indeed intersect.

This shows that $x_1, \dots, x_{d+2}$ cannot be shattered, so $\mathsf{vc}(\mathcal{H}_{\mathsf{lin}}) \le d+1$. Combining this with Proposition 2.24, we obtain:

Proposition 2.30. *With probability $1-\delta$, we have $\widetilde{\mathsf{TV}}_{\mathcal{H}_{\mathsf{lin}}}(\tilde{p}, \tilde{p}_n) \le \mathcal{O}\big(\sqrt{\frac{d+\log(1/\delta)}{n}}\big)$.*

Combining this with Corollary 2.29 and Lemma 2.27, we see that projecting onto $\mathcal{G}_{\mathsf{TV}}(\rho, 2\epsilon')$ under $\widetilde{\mathsf{TV}}_{\mathcal{H}_{\mathsf{lin}}}$ performs well in finite samples, for $\epsilon' = \epsilon + \mathcal{O}(\sqrt{d/n})$. For instance, if $\mathcal{G}$ has bounded covariance we achieve error $\mathcal{O}(\sqrt{\epsilon + \sqrt{d/n}})$; if $\mathcal{G}$ is sub-Gaussian we achieve error $\tilde{\mathcal{O}}(\epsilon + \sqrt{d/n})$; and in general if $\mathcal{G}$ has bounded ψ-norm we achieve error $\mathcal{O}\big((\epsilon + \sqrt{d/n})\psi^{-1}\big(\frac{1}{\epsilon+\sqrt{d/n}}\big)\big) \le \mathcal{O}((\epsilon + \sqrt{d/n})\psi^{-1}(1/\epsilon))$.

This analysis is slightly sub-optimal as the best lower bound we are aware of is $\Omega(\epsilon\psi^{-1}(1/\epsilon) + \sqrt{d/n})$, i.e. the $\psi^{-1}(1/\epsilon)$ coefficient in the dependence on n shouldn't be there. However, it is accurate as long as ϵ is large compared to $\sqrt{d/n}$.

Connection to Tukey median. A classical robust estimator for the mean is the *Tukey median*, which solves the problem

$$\min_{\mu} \max_{v \in \mathbb{R}^d} |\mathbb{P}_{X \sim \tilde{p}_n}[\langle X, v\rangle \ge \langle \mu, v\rangle] - \tfrac{1}{2}| \tag{59}$$

[Note: this definition is slightly wrong as it does not behave gracefully when there is a point mass at μ.]

It is instructive to compare this to projection under $\widetilde{\mathsf{TV}}$, which corresponds to

$$\min_{q \in \mathcal{G}} \max_{v \in \mathbb{R}^d, \tau \in \mathbb{R}} |\mathbb{P}_{X \sim \tilde{p}_n}[\langle X, v\rangle \ge \tau] - \mathbb{P}_{X \sim q}[\langle X, v\rangle \ge \tau]|. \tag{60}$$

The differences are: (1) the Tukey median only minimizes over the mean rather than the full distribution q; (2) it only considers the threshold $\langle \mu, v\rangle$ rather than all thresholds τ; it assumes that the median of any one-dimensional projection $\langle X, v\rangle$ is equal to its mean (which is why we subtract $\frac{1}{2}$ in (59)). Distributions satisfying this final property are said to be *unskewed*.

For unskewed distributions with "sufficient probability mass" near the mean, the Tukey median yields a robust estimator. In fact, it can be robust even if the true distribution has heavy tails (and hence is not resilient), by virtue of leveraging the unskewed property. We will explore this in an exercise.

[Lectures 6-7]

2.6.2 Expanding the Set

In Section 2.6.1 we saw how to resolve the issue with TV projection by relaxing to a weaker distance $\widetilde{\mathsf{TV}}$. We will now study an alternate approach, based on expanding the destination set $\mathcal{G}$ to a larger set $\mathcal{M}$. For this approach we will need to reference the "true empirical distribution" p_n^*. What we mean by this is the following: Whenever $\mathsf{TV}(p^*, \tilde{p}) \le \epsilon$, we know that p^* and $\tilde{p}$ are identical except for some event E of probability ϵ. Therefore we can sample from $\tilde{p}$ as follows:

1. Draw a sample from $X \sim p^*$.
2. Check if E holds; if it does, replace X with a sample from the conditional distribution $\tilde{p}_{|E}$.
3. Otherwise leave X as-is.

Thus we can interpret a sample from $\tilde{p}$ as having a $1-\epsilon$ chance of being "from" p^*. More generally, we can construct the empirical distribution $\tilde{p}_n$ by first constructing the empirical distribution p_n^* coming from p^*, then replacing $\mathsf{Binom}(n, \epsilon)$ of the points with samples from $\tilde{p}_{|E}$. Formally, we have created a coupling between the random variables p_n^* and $\tilde{p}_n$ such that $\mathsf{TV}(p_n^*, \tilde{p}_n)$ is distributed as $\frac{1}{n}\mathsf{Binom}(n, \epsilon)$.

Let us return to expanding the set from $\mathcal{G}$ to $\mathcal{M}$. For this to work, we need three properties to hold:

- $\mathcal{M}$ is large enough: $\min_{q \in \mathcal{M}} \mathsf{TV}(q, p_n^*)$ is small with high probability.
- The empirical loss $L(p_n^*, \theta)$ is a good approximation to the population loss $L(p^*, \theta)$.
- The modulus is still bounded: $\min_{p,q \in \mathcal{M}: \mathsf{TV}(p,q) \le 2\epsilon} L(p, \theta^*(q))$ is small.

In fact, it suffices for $\mathcal{M}$ to satisfy a weaker property; we only need the "generalized modulus" to be small relative to some $\mathcal{G}' \subset \mathcal{M}$:

Proposition 2.31. *For a set $\mathcal{G}' \subset \mathcal{M}$, define the generalized modulus of continuity as*

$$\mathfrak{m}(\mathcal{G}', \mathcal{M}, 2\epsilon) \stackrel{\text{def}}{=} \min_{p \in \mathcal{G}', q \in \mathcal{M}: \mathsf{TV}(p,q) \le 2\epsilon} L(p, \theta^*(q)). \tag{61}$$

Assume that the true empirical distribution p_n^* lies in $\mathcal{G}'$ with probability $1-\delta$. Then the minimum distance functional projecting under TV onto $\mathcal{M}$ has empirical error $L(p_n^*, \hat{\theta})$ at most $\mathfrak{m}(\mathcal{G}', \mathcal{M}, 2\epsilon')$ with probability at least $1 - \delta - \mathbb{P}[\mathsf{Binom}(\epsilon, n) \ge \epsilon' n]$.

Proof. Let $\epsilon' = \mathsf{TV}(p_n^*, \tilde{p}_n)$, which is $\mathsf{Binom}(\epsilon, n)$-distributed. If p_n^* lies in $\mathcal{G}'$, then since $\mathcal{G}' \subset \mathcal{M}$ we know that $\tilde{p}_n$ has distance at most ϵ' from $\mathcal{M}$, and so the projected distribution q satisfies $\mathsf{TV}(q, \tilde{p}_n) \le \epsilon'$ and hence $\mathsf{TV}(q, p_n^*) \le 2\epsilon'$. It follows from the definition that $L(p_n^*, \hat{\theta}) = L(p_n^*, \theta^*(q)) \le \mathfrak{m}(\mathcal{G}', \mathcal{M}, 2\epsilon')$. □

A useful bound on the binomial tail is that $\mathbb{P}[\mathsf{Binom}(\epsilon, n) \ge 2\epsilon n] \le \exp(-\epsilon n/3)$. In particular the empirical error is at most $\mathfrak{m}(\mathcal{G}', \mathcal{M}, 4\epsilon)$ with probability at least $1 - \delta - \exp(-\epsilon n/3)$.

Application: bounded kth moments. First suppose that the distribution p^* has bounded kth moments, i.e. $\mathcal{G}_{\mathsf{mom},k}(\sigma) = \{p \mid \|p\|_\psi \le \sigma\}$, where $\psi(x) = x^k$. When $k > 2$, the empirical distribution p_n^* will not have bouned kth moments until $n \ge \Omega(d^{k/2})$. This is because if we take a single sample $x_1 \sim p$ and let v be a unit vector in the direction of $x_1 - \mu$, then $\mathbb{E}_{x \sim p_n^*}[\langle x - \mu, v\rangle^k] \ge \frac{1}{n}\|x_1 - \mu\|_2^k \gtrsim d^{k/2}/n$, since the norm of $\|x_1 - \mu\|_2$ is typically $\sqrt{d}$.

Consequently, it is necessary to expand the set and we will choose $\mathcal{G}' = \mathcal{M} = \mathcal{G}_{\mathsf{TV}}(\rho, \epsilon)$ for $\rho = \mathcal{O}(\sigma\epsilon^{1-1/k})$ to be the set of resilience distributions with appropriate parameters ρ and ϵ. We already know that the modulus of $\mathcal{M}$ is bounded by $\mathcal{O}(\sigma\epsilon^{1-1/k})$, so the hard part is showing that the empirical distribution p_n^* lies in $\mathcal{M}$ with high probability.

As noted above, we cannot hope to prove that p_n^* has bounded moments except when $n = \Omega(d^{k/2})$, which is too large. We will instead show that certain *truncated* moments of p_n^* are bounded as soon as $n = \Omega(d)$,

and that these truncated moments suffice to show resilience. Specifically, if $\psi(x) = x^k$ is the Orlicz function for the kth moments, we will define the truncated function

$$\tilde{\psi}(x) = \begin{cases} x^k & : \quad x \le x_0 \\ k x_0^{k-1}(x - x_0) + x_0^k & : \quad x > x_0 \end{cases} \tag{62}$$

In other words, $\tilde{\psi}$ is equal to ψ for $x \le x_0$, and is the best linear lower bound to ψ for $x > x_0$. Note that $\tilde{\psi}$ is L-Lipschitz for $L = k x_0^{k-1}$. We will eventually take $x_0 = (k^{k-1}\epsilon)^{-1/k}$ and hence $L = (1/\epsilon)^{(k-1)/k}$. Using a symmetrization argument, we will bound the truncated $\sup_{\|v\|_2 \le 1} \mathbb{E}_{p_n^*}[\tilde{\psi}(|\langle x - \mu, v\rangle|/\sigma)]$.

Proposition 2.32. *Let $X_1, \ldots, X_n \sim p^*$, where $p^* \in \mathcal{G}_{\mathsf{mom},k}(\sigma)$. Then,*

$$\mathbb{E}_{X_1,\ldots,X_n \sim p^*}\left[\left| \sup_{\|v\|_2 \le 1} \frac{1}{n}\sum_{i=1}^n \tilde{\psi}\left(\frac{|\langle X_i - \mu, v\rangle|}{\sigma}\right) - U(v)\right|^k\right] \le O\left(2L\sqrt{\frac{dk}{n}}\right)^k, \tag{63}$$

where $L = k x_0^{k-1}$ and $U(v)$ is a function satisfying $U(v) \le 1$ for all v.

Before proving Proposition 2.32, let us interpret its significance. Take $x_0 = (k^{k-1}\epsilon)^{-1/k}$ and hence $L = \epsilon^{1-1/k}$. Take n large enough so that the right-hand-side of (63) is at most 1, which requires $n \ge \Omega(kd/\epsilon^{2-2/k})$. We then obtain a high-probability bound on the $\tilde{\psi}$-norm of p_n^*, i.e. the $\tilde{\psi}$-norm is at most $\mathcal{O}(\delta^{-1/k})$ with probability $1 - \delta$. This implies that p_n^* is resilient with parameter $\rho = \sigma\epsilon\tilde{\psi}^{-1}(\mathcal{O}(\delta^{-1/k})/\epsilon) = 2\sigma\epsilon^{1-1/k}$. A useful bound on $\tilde{\psi}^{-1}$ is $\tilde{\psi}(-1)(z) \le x_0 + z/L$, and since $x_0 \le (1/\epsilon)^{-1/k}$ and $L = (1/\epsilon)^{(k-1)/k}$ in our case, we have

$$\rho \le \mathcal{O}(\sigma\epsilon^{1-1/k}\delta^{-1/k}) \text{ with probability } 1 - \delta.$$

This matches the population-bound of $\mathcal{O}(\sigma\epsilon^{1-1/k})$, and only requires $kd/\epsilon^{2-2/k}$ samples, in contrast to the d/ϵ^2 samples required before. Indeed, this sample complexity dependence is optimal (other than the factor of k); the only drawback is that we do not get exponential tails (we instead obtain tails of $\delta^{-1/k}$, which is worse than the $\sqrt{\log(1/\delta)}$ from before).

Now we discuss some ideas that are needed in the proof. We would like to somehow exploit the fact that $\tilde{\psi}$ is L-Lipschitz to prove concentration. We can do so with the following keystone result in probability theory:

Theorem 2.33 (Ledoux-Talagrand Contraction)**.** *Let $\phi : \mathbb{R} \to \mathbb{R}$ be an L-Lipschitz function such that $\phi(0) = 0$. Then for any convex, increasing function g and Rademacher variables $\epsilon_{1:n} \sim \{\pm 1\}$, we have*

$$\mathbb{E}_{\epsilon_{1:n}}[g(\sup_{t\in T}\sum_{i=1}^n \epsilon_i \phi(t_i))] \le \mathbb{E}_{\epsilon_{1:n}}[g(L\sup_{t\in T}\sum_{i=1}^n \epsilon_i t_i)]. \tag{64}$$

Let us interpret this result. We should think of the t_i as a quantity such as $\langle x_i - \mu, v\rangle$, where abstracting to t_i yields generality and notational simplicity. Theorem 2.33 says that if we let $Y = \sup_{t\in T}\sum_i \epsilon_i\phi(t_i)$ and $Z = L\sup_{t\in T}\sum_i \epsilon_i t_i$, then $\mathbb{E}[g(Y)] \le \mathbb{E}[g(Z)]$ for all convex increasing functions g. When this holds we say that Y *stochastically dominates Z in second order*; intuitively, it is equivalent to saying that Z has larger mean than Y and greater variation around its mean. For distributions supported on just two points, we can formalize this as follows:

Lemma 2.34 (Two-point stochastic dominance)**.** *Let Y take values y_1 and y_2 with probability $\frac{1}{2}$, and Z take values z_1 and z_2 with probability $\frac{1}{2}$. Then Z stochastically dominates Y (in second order) if and only if*

$$\frac{z_1 + z_2}{2} \ge \frac{y_1 + y_2}{2} \text{ and } \max(z_1, z_2) \ge \max(y_1, y_2). \tag{65}$$

Proof. Without loss of generality assume $z_2 \ge z_1$ and $y_2 \ge y_1$. We want to show that $\mathbb{E}[g(Y)] \le \mathbb{E}[g(Z)]$ for all convex increasing g if and only if (65) holds. We first establish necessity of (65). Take $g(x) = x$, then we require $\mathbb{E}[Y] \le \mathbb{E}[Z]$, which is the first condition in (65). Taking $g(x) = \max(x - z_2, 0)$ yields $\mathbb{E}[g(Z)] = 0$ and $\mathbb{E}[g(Y)] \ge \frac{1}{2}\max(y_2 - z_2, 0)$, so $\mathbb{E}[g(Y)] \le \mathbb{E}[g(Z)]$ implies that $y_2 \le z_2$, which is the second condition in (65).

We next establish sufficiency, by conjuring up appropriate weights for Jensen's inequality. We have

$$\frac{y_2 - z_1}{z_2 - z_1} g(z_2) + \frac{z_2 - y_2}{z_2 - z_1} g(z_1) \geq g\Big(\frac{z_2(y_2 - z_1) + z_1(z_2 - y_2)}{z_2 - z_1}\Big) = g(y_2), \tag{66}$$

$$\frac{z_2 - y_2}{z_2 - z_1} g(z_2) + \frac{y_2 - z_1}{z_2 - z_1} g(z_1) \geq g\Big(\frac{z_2(z_2 - y_2) + z_1(y_2 - z_1)}{z_2 - z_1}\Big) = g(z_1 + z_2 - y_2) \geq g(y_1). \tag{67}$$

Here the first two inequalities are Jensen while the last is by the first condition in (65) together with the monotonicity of g. Adding these together yields $g(z_2) + g(z_1) \geq g(y_2) + g(y_1)$, or $\mathbb{E}[g(Z)] \geq \mathbb{E}[g(Y)]$, as desired. We need only check that the weights $\frac{y_2 - z_1}{z_2 - z_1}$ and $\frac{z_2 - y_2}{z_2 - z_1}$ are positive. The second weight is positive by the assumption $z_2 \geq y_2$. The first weight could be negative if $y_2 < z_1$, meaning that *both* y_1 and y_2 are smaller than *both* z_1 and z_2. But in this case, the inequality $\mathbb{E}[g(Y)] \leq \mathbb{E}[g(Z)]$ trivially holds by monotonicity of g. This completes the proof. □

We are now ready to prove Theorem 2.33.

Proof of Theorem 2.33. Without loss of generality we may take $L = 1$. Our strategy will be to iteratively apply an inequality for a single ϵ_i to replace all the $\phi(t_i)$ with t_i one-by-one. The inequality for a single ϵ_i is the following:

Lemma 2.35. *For any 1-Lipschitz function ϕ with $\phi(0) = 0$, any collection T of ordered pairs (a, b), and any convex increasing function g, we have*

$$\mathbb{E}_{\epsilon \sim \{-1,+1\}}[g(\sup_{(a,b)\in T} a + \epsilon\phi(b))] \leq \mathbb{E}_{\epsilon \sim \{-1,+1\}}[g(\sup_{(a,b)\in T} a + \epsilon b)]. \tag{68}$$

To prove this, let (a_+, b_+) attain the sup of $a + \epsilon\phi(b)$ for $\epsilon = +1$, and (a_-, b_-) attain the sup for $\epsilon = -1$. We will check the conditions of Lemma 2.34 for

$$y_1 = a_- - \phi(b_-), \tag{69}$$

$$y_2 = a_+ + \phi(b_+), \tag{70}$$

$$z_1 = \max(a_- - b_-, a_+ - b_+), \tag{71}$$

$$z_2 = \max(a_- + b_-, a_+ + b_+). \tag{72}$$

(Note that z_1 and z_2 are lower-bounds on the right-hand-side sup for $\epsilon = -1, +1$ respectively.)

First we need $\max(y_1, y_2) \leq \max(z_1, z_2)$. But $\max(z_1, z_2) = \max(a_- + |b_-|, a_+ + |b_+|) \geq \max(a_- - \phi(b_-), a_+ + \phi(b_+)) = \max(y_1, y_2)$. Here the inequality follows since $\phi(b) \leq |b|$ since ϕ is Lipschitz and $\phi(0) = 0$.

Second we need $\frac{y_1+y_2}{2} \leq \frac{z_1+z_2}{2}$. We have $z_1 + z_2 \geq \max((a_- - b_-) + (a_+ + b_+), (a_- + b_-) + (a_+ - b_+)) = a_+ + a_- + |b_+ - b_-|$, so it suffices to show that $\frac{a_+ + a_- + |b_+ - b_-|}{2} \geq \frac{a_+ + a_- + \phi(b_+) - \phi(b_-)}{2}$. This exactly reduces to $\phi(b_+) - \phi(b_-) \leq |b_+ - b_-|$, which again follows since ϕ is Lipschitz. This completes the proof of the lemma.

Now to prove the general proposition we observe that if $g(x)$ is convex in x, so is $g(x + t)$ for any t. We

then proceed by iteratively applying Lemma 2.35:

$$\mathbb{E}_{\epsilon_{1:n}}[g(\sup_{t\in T}\sum_{i=1}^{n}\epsilon_i\phi(t_i))] = \mathbb{E}_{\epsilon_{1:n-1}}[\mathbb{E}_{\epsilon_n}[g(\sup_{t\in T}\underbrace{\sum_{i=1}^{n-1}\epsilon_i\phi(t_i)}_{a} + \epsilon_n\underbrace{\phi(t_n)}_{\phi(b)}) \mid \epsilon_{1:n-1}]] \tag{73}$$

$$\leq \mathbb{E}_{\epsilon_{1:n-1}}[\mathbb{E}_{\epsilon_n}[g(\sup_{t\in T}\sum_{i=1}^{n-1}\epsilon_i\phi(t_i) + \epsilon_n t_n) \mid \epsilon_{1:n-1}]] \tag{74}$$

$$= \mathbb{E}_{\epsilon_{1:n}}[[g(\sup_{t\in T}\sum_{i=1}^{n-1}\epsilon_i\phi(t_i) + \epsilon_n t_n)] \tag{75}$$

$$\vdots \tag{76}$$

$$\leq \mathbb{E}_{\epsilon_{1:n}}[g(\sup_{t\in T}\epsilon_1\phi(t_1) + \sum_{i=2}^{n}\epsilon_i t_i)] \tag{77}$$

$$\leq \mathbb{E}_{\epsilon_{1:n}}[g(\sup_{t\in T}\sum_{i=1}^{n}\epsilon_i t_i)], \tag{78}$$

which completes the proof. □

Let us return now to bounding the truncated moments in Proposition 2.32.

Proof of Proposition 2.32. We start with a symmetrization argument. Let $\mu_{\tilde{\psi}} = \mathbb{E}_{X\sim p^*}[\tilde{\psi}(|\langle X-\mu, v\rangle|/\sigma)]$, and note that $\mu_{\tilde{\psi}} \leq \mu_\psi \leq 1$. Now, by symmetrization we have

$$\mathbb{E}_{X_1,\ldots,X_n\sim p^*}\left[\left|\sup_{\|v\|_2\leq 1}\frac{1}{n}\sum_{i=1}^{n}\tilde{\psi}\Big(\frac{|\langle X_i-\mu, v\rangle|}{\sigma}\Big) - \mu_{\tilde{\psi}}\right|^k\right] \tag{79}$$

$$\leq \mathbb{E}_{X,X'\sim p,\epsilon}\left[\left|\sup_{\|v\|_2\leq 1}\frac{1}{n}\sum_{i=1}^{n}\epsilon_i\Big(\tilde{\psi}\Big(\frac{|\langle X_i-\mu, v\rangle|}{\sigma}\Big) - \tilde{\psi}\Big(\frac{|\langle X_i'-\mu, v\rangle|}{\sigma}\Big)\Big)\right|^k\right] \tag{80}$$

$$\leq 2^k\mathbb{E}_{X\sim p,\epsilon}\left[\left|\sup_{\|v\|_2\leq 1}\frac{1}{n}\sum_{i=1}^{n}\epsilon_i\tilde{\psi}\Big(\frac{|\langle X_i-\mu, v\rangle|}{\sigma}\Big)\right|^k\right]. \tag{81}$$

Here the first inequality adds and subtracts the mean, the second applies symmetrization, while the third uses the fact that optimizing a single v for both X and X' is smaller than optimizing v separately for each (and that the expectations of the expressions with X and X' are equal to each other in that case).

We now apply Ledoux-Talagrand contraction. Invoking Theorem 2.33 with $g(x) = |x|^k$, $\phi(x) = \tilde{\psi}(|x|)$ and $t_i = \langle X_i-\mu, v\rangle|/\sigma$, we obtain

$$\mathbb{E}_{X\sim p,\epsilon}\left[\left|\sup_{\|v\|_2\leq 1}\frac{1}{n}\sum_{i=1}^{n}\epsilon_i\tilde{\psi}\Big(\frac{|\langle X_i-\mu, v\rangle|}{\sigma}\Big)\right|^k\right] \leq (L/\sigma)^k\mathbb{E}_{X\sim p,\epsilon}\Big[\Big|\sup_{\|v\|_2\leq 1}\frac{1}{n}\sum_{i=1}^{n}\epsilon_i\langle X_i-\mu, v\rangle\Big|^k\Big] \tag{82}$$

$$= (L/\sigma)^k\mathbb{E}_{X\sim p,\epsilon}\Big[\Big\|\frac{1}{n}\sum_{i=1}^{n}\epsilon_i(X_i-\mu)\Big\|_2^k\Big]. \tag{83}$$

We are thus finally left to bound $\mathbb{E}_{X\sim p,\epsilon}[\|\sum_{i=1}^n\epsilon_i(X_i-\mu)\|^k]$. Here we will use *Khintchine's inequality*, which says that

$$A_k\|z\|_2 \leq \mathbb{E}_\epsilon[|\sum_i\epsilon_i z_i|^k]^{1/k} \leq B_k\|z\|_2, \tag{84}$$

where A_k is $\Theta(1)$ and B_k is $\Theta(\sqrt{k})$ for $k\geq 1$. Applying this in our case, we obtain

$$\mathbb{E}_{X,\epsilon}[\|\sum_{i=1}^{n}\epsilon_i(X_i-\mu)\|_2^k] \leq O(1)^k\mathbb{E}_{X,\epsilon,\epsilon'}[|\sum_{i=1}^{n}\epsilon_i\langle X_i-\mu, \epsilon'\rangle|^k]. \tag{85}$$

Next apply Rosenthal's inequality (Eq. 21), which yields that

$$\mathbb{E}_{X,\epsilon}[\sum_{i=1}^n \epsilon_i \langle X_i - \mu, \epsilon' \rangle|^k \mid \epsilon'] \le \mathcal{O}(k)^k \sum_{i=1}^n \mathbb{E}_{X,\epsilon}[|\langle X_i - \mu, \epsilon' \rangle|^k \mid \epsilon'] + \mathcal{O}(\sqrt{k})^k (\sum_{i=1}^n \mathbb{E}[\langle X_i - \mu, \epsilon' \rangle|^2])^{k/2} \quad (86)$$

$$\le \mathcal{O}(k)^k \cdot n\sigma^k \|\epsilon'\|_2^k + \mathcal{O}(\sqrt{kn})^k \sigma^k \|\epsilon'\|_2^k \quad (87)$$

$$= \mathcal{O}(\sigma k\sqrt{d})^k n + \mathcal{O}(\sigma\sqrt{kd})^k n^{k/2}, \quad (88)$$

where the last step uses that $\|\epsilon'\|_2 = \sqrt{d}$ and the second-to-last step uses the bounded moments of X. As long as $n \gg k^{k/(k-2)}$ the latter term dominates and hence plugging back into we conclude that

$$\mathbb{E}_{X,\epsilon}[\| \sum_{i=1}^n \epsilon_i (X_i - \mu)\|_2^k]^{1/k} = \mathcal{O}(\sigma\sqrt{kdn}). \quad (89)$$

Thus bounds the symmetrized truncated moments in (82-83) by $O(L\sqrt{kd/n})^k$, and plugging back into (81) completes the proof. □

Application: isotropic Gaussians. Next take $\mathcal{G}_{\mathsf{gauss}}$ to be the family of isotropic Gaussians $\mathcal{N}(\mu, I)$. We saw earlier that the modulus $\mathfrak{m}(\mathcal{G}_{\mathsf{gauss}}, \epsilon)$ was $\mathcal{O}(\epsilon)$ for the mean estimation loss $L(p, \theta) = \|\theta - \mu(p)\|_2$. Thus projecting onto $\mathcal{G}_{\mathsf{gauss}}$ yields error $\mathcal{O}(\epsilon)$ for mean estimation in the limit of infinite samples, but doesn't work for finite samples since the TV distance to $\mathcal{G}_{\mathsf{gauss}}$ will always be 1.

Instead we will project onto the set $\mathcal{G}_{\mathsf{cov}}(\sigma) = \{p \mid \|\mathbb{E}[(X-\mu)(X-\mu)^\top]\| \le \sigma^2\}$, for $\sigma^2 = \mathcal{O}(1 + d/n + \log(1/\delta)/n)$. We already saw in Lemma 2.23 that when p^* is (sub-)Gaussian the empirical distribution p_n^* lies within this set. But the modulus of $\mathcal{G}_{\mathsf{cov}}$ only decays as $\mathcal{O}(\sqrt{\epsilon})$, which is worse than the $\mathcal{O}(\epsilon)$ dependence that we had in infinite samples! How can we resolve this issue?

We will let $\mathcal{G}_{\mathsf{iso}}$ be the family of distributions whose covariance is not only bounded, but close to the identity, and where moreover this holds for all $(1-\epsilon)$-subsets:

$$\mathcal{G}_{\mathsf{iso}}(\sigma_1, \sigma_2) \stackrel{\text{def}}{=} \{p \mid \|\mathbb{E}_r[X-\mu]\|_2 \le \sigma_1 \text{ and } \|\mathbb{E}_r[(X-\mu)(X-\mu)^\top - I\| \le (\sigma_2)^2, \text{ whenever } r \le \frac{p}{1-\epsilon}\}. \quad (90)$$

The following improvement on Lemma 2.23 implies that $p_n^* \in \mathcal{G}_{\mathsf{iso}}(\sigma_1, \sigma_2)$ for $\sigma_1 = \mathcal{O}(\epsilon\sqrt{\log(1/\epsilon)})$ and $\sigma_2 = \mathcal{O}(\sqrt{\epsilon \log(1/\epsilon)})$. [Note: the lemma below is wrong as stated. To be fixed.]

Lemma 2.36. *Suppose that $X_1, \dots, X_n$ are drawn independently from a sub-Gaussian distribution with sub-Gaussian parameter σ, mean 0, and identity covariance. Then, with probability $1-\delta$ we have*

$$\Big\| \frac{1}{|S|} \sum_{i \in S} X_i X_i^\top - I \Big\| \le \mathcal{O}\Big(\sigma^2 \cdot \Big(\epsilon \log(1/\epsilon) + \frac{d + \log(1/\delta)}{n}\Big)\Big), \text{ and} \quad (91)$$

$$\Big\| \frac{1}{|S|} \sum_{i \in S} X_i \Big\|_2 \le \mathcal{O}\Big(\sigma \cdot \Big(\epsilon\sqrt{\log(1/\epsilon)} + \sqrt{\frac{d + \log(1/\delta)}{n}}\Big)\Big) \quad (92)$$

for all subsets $S \subseteq \{1, \dots, n\}$ with $|S| \ge (1-\epsilon)n$. In particular, if $n \gg d/(\epsilon^2 \log(1/\epsilon))$ then $\delta \le \exp(-c\epsilon n \log(1/\epsilon))$ for some constant c.

We will return to the proof of Lemma 2.36 later. For now, note that this means that $p_n^* \in \mathcal{G}'$ for $\mathcal{G}' = \mathcal{G}_{\mathsf{iso}}(\mathcal{O}(\epsilon\sqrt{\log(1/\epsilon)}, \mathcal{O}(\sqrt{\epsilon \log(1/\epsilon)}))$, at least for large enough n. Furthermore, $\mathcal{G}' \subset \mathcal{M}$ for $\mathcal{M} = \mathcal{G}_{\mathsf{cov}}(1 + \mathcal{O}(\epsilon \log(1/\epsilon)))$.

Now we bound the generalized modulus of continuity:

Lemma 2.37. *Suppose that $p \in \mathcal{G}_{\mathsf{iso}}(\sigma_1, \sigma_2)$ and $q \in \mathcal{G}_{\mathsf{cov}}(\sqrt{1+\sigma_2^2})$, and furthermore $\mathsf{TV}(p, q) \le \epsilon$. Then $\|\mu(p) - \mu(q)\|_2 \le \mathcal{O}(\sigma_1 + \sigma_2\sqrt{\epsilon} + \epsilon)$.*

Proof. Take the midpoint distribution $r = \frac{\min(p,q)}{1-\epsilon}$, and write $q = (1-\epsilon)r + \epsilon q'$. We will bound $\|\mu(r) - \mu(q)\|_2$ (note that $\|\mu(r) - \mu(p)\|_2$ is already bounded since $p \in \mathcal{G}_{\mathsf{iso}}$). We have that

$$\mathsf{Cov}_q[X] = (1-\epsilon)\mathbb{E}_r[(X-\mu_q)(X-\mu_q)^\top] + \epsilon\mathbb{E}_{q'}[(X-\mu_q)(X-\mu_q)^\top] \tag{93}$$

$$= (1-\epsilon)(\mathsf{Cov}_r[X] + (\mu_q-\mu_r)(\mu_q-\mu_r)^\top) + \epsilon\mathbb{E}_{q'}[(X-\mu_q)(X-\mu)q)^\top] \tag{94}$$

$$\succeq (1-\epsilon)(\mathsf{Cov}_r[X] + (\mu_q-\mu_r)(\mu_q-\mu_r)^\top) + \epsilon(\mu_q-\mu_{q'})(\mu_q-\mu_{q'})^\top. \tag{95}$$

A computation yields $\mu_q - \mu_{q'} = \frac{(1-\epsilon)^2}{\epsilon}(\mu_q - \mu_r)$. Plugging this into (95) and simplifying, we obtain that

$$\mathsf{Cov}_q[X] \succeq (1-\epsilon)(\mathsf{Cov}_r[X] + (1/\epsilon)(\mu_q-\mu_r)(\mu_q-\mu_r)^\top). \tag{96}$$

Now since $\mathsf{Cov}_r[X] \succeq (1-\sigma_2^2)I$, we have $\|\mathsf{Cov}_q[X]\| \geq (1-\epsilon)(1-\sigma_2^2) + (1/\epsilon)\|\mu_q - \mu_r\|_2^2$. But by assumption $\|\mathsf{Cov}_q[X]\| \leq 1+\sigma_2^2$. Combining these yields that $\|\mu_r - \mu_q\|_2^2 \leq \epsilon(2\sigma_2^2 + \epsilon + \epsilon\sigma_2^2)$, and so $\|\mu_r - \mu_q\|_2 \leq \mathcal{O}(\epsilon + \sigma_2\sqrt{\epsilon})$, which gives the desired result. □

In conclusion, projecting onto $\mathcal{G}_{\mathsf{cov}}(1 + \mathcal{O}(\epsilon\log(1/\epsilon)))$ under TV distance gives a robust mean estimator for isotropic Gaussians, which achieves error $\mathcal{O}(\epsilon\sqrt{\log(1/\epsilon)})$. This is slightly worse than the optimal $\mathcal{O}(\epsilon)$ bound but improves over the naïve analysis that only gave $\mathcal{O}(\sqrt{\epsilon})$.

Another advantage of projecting onto $\mathcal{G}_{\mathsf{cov}}$ is that, as we will see in Section 2.7, this projection can be done computationally efficiently.

Proof of Lemma 2.36. TBD

[Lecture 8]

2.7 Efficient Algorithms

We now turn our attention to efficient algorithms. Recall that previously we considered minimum distance functionals projecting onto sets $\mathcal{G}$ and $\mathcal{M}$ under distances TV and $\widetilde{\mathsf{TV}}$. Here we will show how to approximately project onto the set $\mathcal{G}_{\mathsf{cov}}(\sigma)$, the family of bounded covariance distributions, under TV distance. The basic idea is that if the true distribution p^* has bounded covariance, and $\tilde{p}$ does not, the largest eigenvector of $\mathsf{Cov}_{\tilde{p}}[X]$ must be well-aligned with the mean of the bad points, and thus we can use this to remove the bad points. If on the other hand $\tilde{p}$ has bounded covariance, then its mean must be close to p^* by our previous modulus bounds and so we are already done.

To study efficient computation we need a way of representing the distributions $\tilde{p}$ and p^*. To do this we will suppose that $\tilde{p}$ is the empirical distribution over n points $x_1, \ldots, x_n$, while p^* is the empirical distribution over some subset S of these points with $|S| \geq (1-\epsilon)n$. Thus in particular p^* is an ϵ-deletion of $\tilde{p}$.

Before we assumed that $\mathsf{TV}(p^*, \tilde{p}) \leq \epsilon$, but taking $p' = \frac{\min(p^*,\tilde{p})}{1-\mathsf{TV}(p^*,\tilde{p})}$, we have $p' \leq \frac{\tilde{p}}{1-\epsilon}$ and $\|\mathsf{Cov}_{p'}[X]\| \leq \frac{\sigma^2}{1-\epsilon} \leq 2\sigma^2$ whenever $\|\mathsf{Cov}_{p^*}[X]\| \leq \sigma^2$. Therefore, taking $p^* \leq \frac{\tilde{p}}{1-\epsilon}$ is equivalent to the TV corruption model from before for our present purposes.

We will construct an efficient algorithm that, given $\tilde{p}$, outputs a distribution q such that $\mathsf{TV}(q, p^*) \leq \mathcal{O}(\epsilon)$ and $\|\mathsf{Cov}_q[X]\|_2 \leq \mathcal{O}(\sigma^2)$. This is similar to the minimum distance functional, in that it finds a distribution close to p^* with bounded covariance; the main difference is that q need not be the projection of $\tilde{p}$ onto $\mathcal{G}_{\mathsf{cov}}$, and also the covariance of q is bounded by $\mathcal{O}(\sigma^2)$ instead of σ^2. However, the modulus of continuity bound from before says that *any* distribution q that is near p^* and has bounded covariance will approximate the mean of p^*. Specifically, we have

$$\|\mu(q) - \mu(p^*)\|_2^2 \leq \mathcal{O}(\max(\|\mathsf{Cov}_q[X]\|, \|\mathsf{Cov}_{p^*}[X]\|) \cdot \mathsf{TV}(p^*, q)) = \mathcal{O}(\sigma^2\epsilon). \tag{97}$$

We will show the following:

Proposition 2.38. *Suppose $\tilde{p}$ and p^* are empirical distributions as above with $p^* \leq \tilde{p}/(1-\epsilon)$, and further suppose that $\|\mathsf{Cov}_{p^*}[X]\| \leq \sigma^2$. Then given $\tilde{p}$ (but not p^*), there is an algorithm with runtime $\mathrm{poly}(n, d)$ that outputs a q with $\mathsf{TV}(p^*, q) \leq \epsilon$ and $\|\mathsf{Cov}_q[X]\| \leq O(\sigma^2)$. In particular, $\|\mu(p^*) - \mu(q)\|_2 = \mathcal{O}(\sigma\sqrt{\epsilon})$.*

Note that the conclusion $\|\mu(p^*) - \mu(q)\|_2 \le \mathcal{O}(\sigma\sqrt{\epsilon})$ follows from the modulus bound on $\mathcal{G}_{\mathsf{cov}}(\sigma)$ together with the property $\mathsf{TV}(p^*, q) \le \epsilon$.

The algorithm, `FilterL2`, underlying Proposition 2.38 is given below; it maintains a weighted distribution $q(c)$, which places weight $c_i / \sum_{j=1}^n c_j$ on point x_i. It then computes the weighted mean and covariance, projects onto the top eigenvector, and downweights points with large projection.

Algorithm 2 `FilterL2`

1: Input: $x_1, \dots, x_n \in \mathbb{R}^d$.
2: Initialize weights $c_1, \dots, c_n = 1$.
3: Compute the empirical mean $\hat{\mu}_c$ of the data, $\hat{\mu}_c \stackrel{\text{def}}{=} (\sum_{i=1}^n c_i x_i)/(\sum_{i=1}^n c_i)$.
4: Compute the empirical covariance $\hat{\Sigma}_c \stackrel{\text{def}}{=} \sum_{i=1}^n c_i (x_i - \hat{\mu}_c)(x_i - \hat{\mu}_c)^\top / \sum_{i=1}^n c_i$.
5: Let v be the maximum eigenvector of $\hat{\Sigma}_c$, and let $\hat{\sigma}_c^2 = v^\top \hat{\Sigma}_c v$.
6: If $\hat{\sigma}_c^2 \le 20\sigma^2$, output $q(c)$.
7: Otherwise, let $\tau_i = \langle x_i - \hat{\mu}_c, v \rangle^2$, and update $c_i \leftarrow c_i \cdot (1 - \tau_i / \tau_{\max})$, where $\tau_{\max} = \max_i \tau_i$.
8: Go back to line 3.

The factor $\tau_{\max}$ in the update $c_i \leftarrow c_i \cdot (1 - \tau_i/\tau_{\max})$ is so that the weights remain positive; the specific factor is unimportant and the main property required is that each point is downweighted proportionally to τ_i. Note also that Algorithm 2 must eventually terminate because one additional weight c_i is set to zero in every iteration of the algorithm.

The intuition behind Algorithm 2 is as follows: if the empirical variance $\hat{\sigma}_c^2$ is much larger than the variance σ^2 of the good data, then the bad points must on average be very far away from the empirical mean (i.e., τ_i must be large on average for the bad points).

More specifically, note that $\tau_i = \langle x_i - \hat{\mu}_c, v^* \rangle^2$. Let $\tilde{\tau}_i = \langle x_i - \mu, v^* \rangle^2$, and imagine for now that $\tau_i \approx \tilde{\tau}_i$. We know that the average of $\tilde{\tau}_i$ over the good points is at most σ^2, since $\tilde{\tau}_i$ is the variance along the projection v^* and $\|\mathsf{Cov}_{p^*}[X]\| \le \sigma^2$. Thus if the overall average of the τ_i is large (say $20\sigma^2$), it must be on account of the bad points. But since there are not that many bad points, their average must be *quite* large–on the order of σ^2/ϵ. Thus they should be easy to separate from the good points. This is depicted in Figure 7.

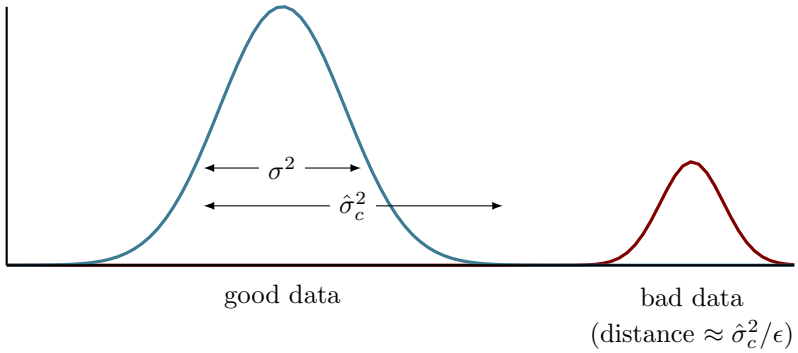

Figure 7: Intuition behind Algorithm 2. Because there is only an ϵ-fraction of bad data, it must lie far away to increase the variance by a constant factor.

This is the basic idea behind the proof, but there are a couple issues with this:

- The assumption that $\tilde{\tau}_i \approx \tau_i$ is basically an assumption that $\mu \approx \hat{\mu}_c$ (which is what we are trying to show in the first place!).
- The bad points are not deterministically larger than the good points; they are only separated in expected value.
- There are many fewer bad points than good points, so they are harder to find.

We will deal with the first issue by showing that μ is close enough to $\hat{\mu}_c$ for the algorithm to make progress. The second issue is why we need to do soft downweighting rather than picking a hard threshold and removing all points with τ_i above the threshold. We will resolve the third issue by showing that we always remove more mass c_i from the bad points than from the good points when we update c_i. Intuitively, while there are only ϵ times as many bad points as good points, this is balanced against the fact that the mean of the bad points is $1/\epsilon$ times as large as the mean of the good points.

We next put this intuition together into a formal proof.

Proof of Proposition 2.38. As above, for weights $c_i \in [0,1]$, let $q(c)$ be the distribution that assigns weight $c_i / \sum_j c_j$ to point x_i. Thus when $c_i = 1$ for all i, we have $q(c) = \tilde{p}$. Our hope is that as the algorithm progresses $q(c)$ approaches p^* or at least has small covariance. We will establish the following invariant:

$$\mathsf{TV}(q(c), p^*) \le \frac{\epsilon}{1-\epsilon} \text{ for all weight vectors } c \text{ used during the execution of Algorithm 2.} \tag{$\mathcal{I}_1$}$$

We will do this by proving the following more complex invariant, which we will later show implies ($\mathcal{I}_1$):

$$\sum_{i \in S}(1-c_i) \le \sum_{i \notin S}(1-c_i) \tag{$\mathcal{I}_2$}$$

The invariant ($\mathcal{I}_2$) says that the total probability mass removed from the good points is less than the total probability mass removed from the bad points. A key lemma relates ($\mathcal{I}_2$) to the τ_i:

Lemma 2.39. *If ($\mathcal{I}_2$) and $\sum_{i \in S} c_i \tau_i \le \sum_{i \notin S} c_i \tau_i$, then it continues to hold after the update $c_i' = c_i(1 - \tau_i/\tau_{\max})$.*

Proof. For any set T, we have

$$\sum_{i \in T} 1 - c_i' = \sum_{i \in T}(1-c_i) + \sum_{i \in T}(c_i - c_i') = \sum_{i \in T}(1-c_i) + \frac{1}{\tau_{\max}} \sum_{i \in T} c_i \tau_i. \tag{98}$$

Applying this for $T = S$ and $T = [n] \backslash S$ yields the lemma. □

Thus our main job is to show that $\sum_{i \in S} c_i \tau_i \le \sum_{i \notin S} c_i \tau_i$. Equivalently, we wish to show that $\sum_{i \in S} c_i \tau_i \le \frac{1}{2} \sum_{i=1}^n c_i \tau_i$. For this, the following bound is helpful:

$$\begin{aligned}
\sum_{i \in S} c_i \tau_i &= \sum_{i \in S} c_i \langle x_i - \hat{\mu}_c, v^* \rangle^2 && (99)\\
&\le \sum_{i \in S} \langle x_i - \hat{\mu}_c, v^* \rangle^2 && (100)\\
&= (1-\epsilon) n \mathbb{E}_{p^*}[\langle x_i - \hat{\mu}_c, v^* \rangle^2] && (101)\\
&= (1-\epsilon) n \cdot (v^*)^\top (\mathsf{Cov}_{p^*}[X] + (\mu - \hat{\mu}_c)(\mu - \hat{\mu}_c)^\top)(v^*) && (102)\\
&\le (1-\epsilon) n \cdot (\|\mathsf{Cov}_{p^*}[X]\| + \|\mu - \hat{\mu}_c\|_2^2). && (103)
\end{aligned}$$

Here the second-to-last step uses the fact that for any θ, $\mathbb{E}[(X-\theta)(X-\theta)^\top] = \mathsf{Cov}[X] + (\theta - \mu)(\theta - \mu)^\top$.

Next note that $\|\mathsf{Cov}_{p^*}\| \le \sigma^2$ while $\|\mu - \hat{\mu}_c\|_2^2 \le \frac{8\epsilon}{1-2\epsilon}\sigma_c^2$ by the modulus of continuity bound combined with the fact that $p^*, q(c) \in \mathcal{G}_{\mathsf{cov}}(\hat{\sigma})$ and $\mathsf{TV}(p^*, q(c)) \le \frac{\epsilon}{1-\epsilon}$. Therefore, we have

$$\sum_{i \in S} c_i \tau_i \le (1-\epsilon)\sigma^2 n + \frac{8\epsilon(1-\epsilon)}{1-2\epsilon}\hat{\sigma}_c^2 n. \tag{104}$$

On the other hand, we have

$$\sum_{i=1}^n c_i \tau_i = (\sum_{i=1}^n c_i) \|\mathsf{Cov}_{q(c)}[X]\| = (\sum_{i=1}^n c_i)\hat{\sigma}_c^2 \ge (1-2\epsilon)\hat{\sigma}_c^2 n, \tag{105}$$

where the final inequality uses the fact that we have so far removed more mass from bad points than good points and hence at most 2ϵ mass in total. Recalling that we wish to show that (104) is at most half of (105), we require that

$$(1-2\epsilon)\hat{\sigma}_c^2 \geq 2(1-\epsilon)\sigma^2 + \frac{16\epsilon(1-\epsilon)}{1-2\epsilon}\hat{\sigma}_c^2, \tag{106}$$

which upon re-arrangement yields

$$\hat{\sigma}_c^2 \geq \frac{2(1-\epsilon)(1-2\epsilon)}{1-12\epsilon+12\epsilon^2}\sigma^2 \tag{107}$$

Since $\hat{\sigma}_c^2 \geq 20\sigma^2$ whenever the algorithm does not terminate, this holds as long as $\epsilon \leq \frac{1}{12}$ (then the constant in front of σ^2 is $\frac{55}{3} < 20$). This shows that ($\mathcal{I}_2$) holds throughout the algorithm.

The one remaining detail is to prove that ($\mathcal{I}_2$) implies ($\mathcal{I}_1$). We wish to show that $\mathsf{TV}(p^*, q(c)) \leq \frac{\epsilon}{1-\epsilon}$. We use the following formula for TV: $\mathsf{TV}(p,q) = \int \max(q(x) - p(x), 0)dx$. Let β be such that $\sum_{i=1}^n c_i = (1-\beta)n$. Then we have

$$\mathsf{TV}(p^*, q(c)) = \sum_{i \in S} \max\left(\frac{c_i}{(1-\beta)n} - \frac{1}{(1-\epsilon)n}, 0\right) + \sum_{i \notin S} \frac{c_i}{(1-\beta)n}. \tag{108}$$

If $\beta \leq \epsilon$, then the first sum is zero while the second sum is at most $\frac{\epsilon}{1-\beta} \leq \frac{\epsilon}{1-\epsilon}$. If on the other hand $\beta > \epsilon$, we will instead use the equality obtained by swapping p and q, which yields

$$\mathsf{TV}(p^*, q(c)) = \sum_{i \in S} \max\left(\frac{1}{(1-\epsilon)n} - \frac{c_i}{(1-\beta)n}, 0\right) \tag{109}$$

$$= \frac{1}{(1-\epsilon)(1-\beta)n} \sum_{i \in S} \max((1-\beta)(1-c_i) + (\epsilon-\beta)c_i, 0). \tag{110}$$

Since $(\epsilon - \beta)c_i \leq 0$ and $\sum_{i \in S}(1-c_i) \leq \epsilon n$, this yields a bound of $\frac{(1-\beta)\epsilon}{(1-\epsilon)(1-\beta)} = \frac{\epsilon}{1-\epsilon}$. We thus obtain the desired bound no matter the value of β, so $\mathsf{TV}(p^*, q(c)) \leq \frac{\epsilon}{1-\epsilon}$ whenever ($\mathcal{I}_2$) holds. This completes the proof. □

[Lecture 9]

2.7.1 Approximate Eigenvectors in Other Norms

Algorithm 2 is specific to the ℓ_2-norm. Let us suppose that we care about recovering an estimate $\hat{\mu}$ such that $\|\mu - \hat{\mu}\|$ is small in some norm other than ℓ_2 (such as the ℓ_1-norm, which may be more appropriate for some combinatorial problems). It turns out that an analog of bounded covariance is sufficient to enable estimation with the typical $\mathcal{O}(\sigma\sqrt{\epsilon})$ error, as long as we can approximately solve the analogous eigenvector problem. To formalize this, we will make use of the *dual norm*:

Definition 2.40. Given a norm $\|\cdot\|$, the *dual norm* $\|\cdot\|_*$ is defined as

$$\|u\|_* = \sup_{\|v\|_2 \leq 1} \langle u, v \rangle. \tag{111}$$

As some examples, the dual of the ℓ_2-norm is itself, the dual of the ℓ_1-norm is the ℓ_∞-norm, and the dual of the ℓ_∞-norm is the ℓ_1-norm. An important property (we omit the proof) is that the dual of the dual is the original norm:

Proposition 2.41. *If $\|\cdot\|$ is a norm on a finite-dimensional vector space, then $\|\cdot\|_{**} = \|\cdot\|$.*

For a more complex example: let $\|v\|_{(k)}$ be the sum of the k largest coordinates of v (in absolute value). Then the dual of $\|\cdot\|_{(k)}$ is $\max(\|u\|_\infty, \|u\|_1/k)$. This can be seen by noting that the vertices of the constraint set $\{u \mid \|u\|_\infty \leq 1, \|u\|_1 \leq k\}$ are exactly the k-sparse $\{-1, 0, +1\}$-vectors.

Let $\mathcal{G}_{\mathsf{cov}}(\sigma, \|\cdot\|)$ denote the family of distributions satisfying $\max_{\|v\|_* \leq 1} v^\top \mathsf{Cov}_p[X]v \leq \sigma^2$. Then $\mathcal{G}_{\mathsf{cov}}$ is resilient exactly analogously to the ℓ_2-case:

Proposition 2.42. *If $p \in \mathcal{G}_{\mathsf{cov}}(\sigma, \|\cdot\|)$ and $r \le \frac{p}{1-\epsilon}$, then $\|\mu(r) - \mu(p)\| \le \sqrt{\frac{2\epsilon}{1-\epsilon}}\sigma$. In other words, all distributions in $\mathcal{G}_{\mathsf{cov}}(\sigma, \|\cdot\|)$ are $(\epsilon, \mathcal{O}(\sigma\sqrt{\epsilon}))$-resilient.*

Proof. We have that $\|\mu(r) - \mu(p)\| = \langle \mu(r) - \mu(p), v\rangle$ for some vector v with $\|v\|_* = 1$. The result then follows by resilience for the one-dimensional distribution $\langle X, v\rangle$ for $X \sim p$. □

When $p^* \in \mathcal{G}_{\mathsf{cov}}(\sigma, \|\cdot\|)$, we will design efficient algorithms analogous to Algorithm 2. The main difficulty is that in norms other than ℓ_2, it is generally not possible to exactly solve the optimization problem $\max_{\|v\|_* \le 1} v^\top \hat{\Sigma}_c v$ that is used in Algorithm 2. We instead make use of a *κ-approximate oracle*:

Definition 2.43. A function $\mathcal{A}(\Sigma)$ is a κ-approximate oracle if for all Σ, $M = \mathcal{A}(\Sigma)$ is a positive semidefinite matrix satisfying

$$\langle M, \Sigma\rangle \ge \sup_{\|v\|_* \le 1} v^\top \Sigma v, \text{ and } \langle M, \Sigma'\rangle \le \kappa \sup_{\|v\|_* \le 1} v^\top \Sigma' v \text{ for all } \Sigma' \succeq 0. \tag{112}$$

Thus a κ-approximate oracle over-approximates $\langle vv^\top, \Sigma$ for the maximizing vector v on Σ, and it underapproximates $\langle vv^\top, \Sigma'\rangle$ within a factor of κ for all $\Sigma' \ne \Sigma$. Given such an oracle, we have the following analog to Algorithm 2:

Algorithm 3 `FilterNorm`

1: Initialize weights $c_1, \dots, c_n = 1$.
2: Compute the empirical mean $\hat{\mu}_c$ of the data, $\hat{\mu}_c \stackrel{\text{def}}{=} (\sum_{i=1}^n c_i x_i)/(\sum_{i=1}^n c_i)$.
3: Compute the empirical covariance $\hat{\Sigma}_c \stackrel{\text{def}}{=} \sum_{i=1}^n c_i (x_i - \hat{\mu}_c)(x_i - \hat{\mu}_c)^\top / \sum_{i=1}^n c_i$.
4: Let $M = \mathcal{A}(\hat{\Sigma}_c)$ be the output of a κ-approximate oracle.
5: If $\langle M, \hat{\Sigma}_c\rangle \le 20\kappa\sigma^2$, output $q(c)$.
6: Otherwise, let $\tau_i = (x_i - \hat{\mu}_c)^\top M (x_i - \hat{\mu}_c)$, and update $c_i \leftarrow c_i \cdot (1 - \tau_i/\tau_{\max})$, where $\tau_{\max} = \max_i \tau_i$.
7: Go back to line 2.

Algorithm 3 outputs an estimate of the mean with error $\mathcal{O}(\sigma\sqrt{\kappa\epsilon})$. The proof is almost exactly the same as Algorithm 2; the main difference is that we need to ensure that $\langle \Sigma, M\rangle$, the inner product of M with the true covariance, is not too large. This is where we use the κ-approximation property. We leave the detailed proof as an exercise, and focus on how to construct a κ-approximate oracle $\mathcal{A}$.

Semidefinite programming. As a concrete example, suppose that we wish to estimate μ in the ℓ_1-norm $\|v\| = \sum_{j=1}^d |v_j|$. The dual norm is the ℓ_∞-norm, and hence our goal is to approximately solve the optimization problem

$$\text{maximize } v^\top \Sigma v \text{ subject to } \|v\|_\infty \le 1. \tag{113}$$

The issue with (113) is that it is not concave in v because of the quadratic function $v^\top \Sigma v$. However, note that $v^\top \Sigma v = \langle \Sigma, vv^\top\rangle$. Therefore, if we replace v with the variable $M = vv^\top$, then we can re-express the optimization problem as

$$\text{maximize } \langle \Sigma, M\rangle \text{ subject to } M_{jj} \le 1 \text{ for all j, } M \succeq 0, \text{rank}(M) = 1. \tag{114}$$

Here the first constraint is a translation of $\|v\|_\infty \le 1$, while the latter two constrain M to be of the form $vv^\top$.

This is almost convex in M, except for the constraint $\text{rank}(M) = 1$. If we omit this constraint, we obtain the optimization

$$\begin{aligned} &\text{maximize } \langle \Sigma, M\rangle \\ &\text{subject to } M_{jj} = 1 \text{ for all } j, \\ &\qquad M \succeq 0. \end{aligned} \tag{115}$$

Note that here we replace the constraint $M_{jj} \le 1$ with $M_{jj} = 1$; this can be done because the maximizer of (115) will always have $M_{jj} = 1$ for all j. For brevity we often write this constraint as $\text{diag}(M) = 1$.

The problem (115) is a special instance of a *semidefinite program* and can be solved in polynomial time (in general, a semidefinite program allows arbitrary linear inequality or positive semidefinite constraints between linear functions of the decision variables; we discuss this more below).

The optimizer M^* of (115) will always satisfy $\langle \Sigma, M^* \rangle \geq \sup_{\|v\|_\infty \leq 1} v^\top \Sigma v$ because and v with $\|v\|_\infty \leq 1$ yields a feasible M. The key is to show that it is not too much larger than this. This turns out to be a fundamental fact in the theory of optimization called *Grothendieck's inequality*:

Theorem 2.44. *If $\Sigma \succeq 0$, then the value of (115) is at most $\frac{\pi}{2} \sup_{\|v\|_\infty \leq 1} v^\top \Sigma v$.*

See Alon and Naor (2004) for a very well-written exposition on Grothendieck's inequality and its relation to optimization algorithms. In that text we also see that a version of Theorem 2.44 holds even when Σ is not positive semidefinite or indeed even square. Here we produce a proof based on [todo: cite] for the semidefinite case.

Proof of Theorem 2.44. The proof involves two key relations. To describe the first, given a matrix X let $\arcsin[X]$ denote the matrix whose i, j entry is $\arcsin(X_{ij})$ (i.e. we apply arcsin element-wise). Then we have (we will show this later)

$$\max_{\|v\|_\infty \leq 1} v^\top \Sigma v = \max_{X \succeq 0, \mathrm{diag}(X)=1} \frac{2}{\pi} \langle \Sigma, \arcsin[X] \rangle. \tag{116}$$

The next relation is that

$$\arcsin[X] \succeq X. \tag{117}$$

Together, these imply the approximation ratio, because we then have

$$\max_{M \succeq 0, \mathrm{diag}(M)=1} \langle \Sigma, M \rangle \leq \max_{M \succeq 0, \mathrm{diag}(M)=1} \langle \Sigma, \arcsin[M] \rangle = \frac{\pi}{2} \max_{\|v\|_\infty \leq 1} v^\top \Sigma v. \tag{118}$$

We will therefore focus on establishing (116) and (117).

To establish (116), we will show that any X with $X \succeq 0$, $\mathrm{diag}(X) = 1$ can be used to produce a probability distribution over vectors v such that $\mathbb{E}[v^\top \Sigma v] = \frac{2}{\pi} \langle \Sigma, \arcsin[X] \rangle$.

First, by Graham/Cholesky decomposition we know that there exist vectors u_i such that $M_{ij} = \langle u_i, u_j \rangle$ for all i, j. In particular, $M_{ii} = 1$ implies that the u_i have unit norm. We will then construct the vector v by taking $v_i = \mathrm{sign}(\langle u_i, g \rangle)$ for a Gaussian random variable $g \sim \mathcal{N}(0, I)$.

We want to show that $\mathbb{E}_g[v_i v_j] = \frac{2}{\pi} \arcsin(\langle u_i, u_j \rangle)$. For this it helps to reason in the two-dimensional space spanned by v_i and v_j. Then $v_i v_j = -1$ if the hyperplane induced by g cuts between u_i and u_j, and $+1$ if it does not. Letting θ be the angle between u_i and u_j, we then have $\mathbb{P}[v_j v_j = -1] = \frac{\theta}{\pi}$ and hence

$$\mathbb{E}_g[v_i v_j] = (1 - \frac{\theta}{\pi}) - \frac{\theta}{\pi} = \frac{2}{\pi}(\frac{\pi}{2} - \theta) = \frac{2}{\pi} \arcsin(\langle u_i, u_j \rangle), \tag{119}$$

as desired. Therefore, we can always construct a distribution over v for which $\mathbb{E}[v^\top \Sigma v] = \frac{2}{\pi} \langle \Sigma, \arcsin[M] \rangle$, hence the right-hand-side of (116) is at most the left-hand-side. For the other direction, note that the maximizing v on the left-hand-side is always a $\{-1, +1\}$ vector by convexity of $v^\top \Sigma v$, and for any such vector we have $\frac{2}{\pi} \arcsin[vv^\top] = vv^\top$. Thus the left-hand-side is at most the right-hand-side, and so the equality (116) indeed holds.

We now turn our attention to establishing (117). For this, let $X^{\odot k}$ denote the matrix whose i, j entry is X_{ij}^k (we take element-wise power). We require the following lemma:

Lemma 2.45. *For all $k \in \{1, 2, \ldots\}$, if $X \succeq 0$ then $X^{\odot k} \succeq 0$.*

Proof. The matrix $X^{\odot k}$ is a submatrix of $X^{\otimes k}$, where $(X^{\otimes k})_{i_1 \cdots i_k, j_1 \cdots j_k} = X_{i_1, j_1} \cdots X_{i_k, j_k}$. We can verify that $X^{\otimes k} \succeq 0$ (its eigenvalues are $\lambda_{i_1} \cdots \lambda_{i_k}$ where λ_i are the eigenvalues of X), hence so is $X^{\odot k}$ since submatrices of PSD matrices are PSD. $\square$

With this in hand, we also make use of the Taylor series for $\arcsin(z)$: $\arcsin(z) = \sum_{n=0}^{\infty} \frac{(2n)!}{(2^n n!)^2} \frac{z^{2n+1}}{2n+1} = z + \frac{z^3}{6} + \cdots$. Then we have

$$\arcsin[X] = X + \sum_{n=1}^{\infty} \frac{(2n)!}{(2^n n!)^2} \frac{1}{2n+1} X^{\odot(2n+1)} \succeq X, \tag{120}$$

as was to be shown. This completes the proof. □

Alternate proof (by Mihaela Curmei): We can also show that $X^{\odot k} \succeq 0$ more directly. Specifically, we will show that if $A, B \succeq 0$ then $A \odot B \succeq 0$, from which the result follows by induction. To show this let $A = \sum_i \lambda_i u_i u_i^\top$ and $B = \sum_j \nu_j v_j v_j^\top$ and observe that

$$A \odot B = (\sum_i \lambda_i u_i u_i^\top) \odot (\sum_j \nu_j v_j v_j^\top) \tag{121}$$

$$= \sum_{i,j} \lambda_i \nu_j (u_i u_i^\top) \odot (v_j v_j^\top) \tag{122}$$

$$= \sum_{i,j} \underbrace{\lambda_i \nu_j}_{\geq 0} \underbrace{(u_i \odot v_j)(u_i \odot v_j)^\top}_{\succeq 0}, \tag{123}$$

from which the claim follows. Here the key step is that for rank-one matrices the $\odot$ operation behaves nicely: $(u_i u_i^\top) \odot (v_j v_j^\top) = (u_i \odot v_j)(u_i \odot v_j)^\top$.

[Lecture 10]

2.8 Semidefinite Programming and Sum-of-Squares

In the previous subsection, we saw how to approximately solve $\max_{\|v\|_\infty \leq 1} v^\top \Sigma v$ via the semidefinite program defined by $\max_{M \succeq 0, \mathrm{diag}(M)=1} \langle M, \Sigma \rangle$. In this section we will cover semidefinite programming in more detail, and build up to *sum-of-squares programming*, which will be used to achieve error $\mathcal{O}(\epsilon^{1-1/k})$ when p^* has "certifiably bounded" kth moments (recall that we earlier achieved error $\mathcal{O}(\epsilon^{1-1/k})$ for bounded kth moments but did not have an efficient algorithm).

A **semidefinite program** is an optimization problem of the form

$$\begin{aligned} \text{maximize } & \langle A, X \rangle \\ \text{subject to } & X \succeq 0, \\ & \langle X, B_1 \rangle \leq c_1, \\ & \vdots \\ & \langle X, B_m \rangle \leq c_m. \end{aligned} \tag{124}$$

Here $\langle X, Y \rangle = \mathrm{tr}(X^T Y) = \sum_{ij} X_{ij} Y_{ij}$ is the inner product between matrices, which is the same as the elementwise dot product when considered as n^2-dimensional vectors.

Here the matrix A specifies the objective of the program, while (B_j, c_j) specify linear inequality constraints. We additionally have the positive semidefinite cone constraint that $X \succeq 0$, meaning that X must be symmetric with only non-negative eigenvalues. Each of A and $B_1, \ldots, B_m$ are $n \times n$ matrices while the c_j are scalars. We can equally well minimize as maximize by replacing A with $-A$.

While (124) is the canonical form for a semidefinite program, problems that are seemingly more complex can be reduced to this form. For one, we can add linear equality constraints as two-sided inequality constraints. In addition, we can replace $X \succeq 0$ with $\mathcal{L}(X) \succeq 0$ for any linear function $\mathcal{L}$, by using linear equality constraints to enforce the linear relations implied by $\mathcal{L}$. Finally, we can actually include any number of constraints $\mathcal{L}_1(X) \succeq 0$, $\mathcal{L}_k(X) \succeq 0$, since this is e.g. equivalent to the single constraint $\begin{bmatrix} \mathcal{L}_1(X) & 0 \\ 0 & \mathcal{L}_2(X) \end{bmatrix}$ when

$k = 2$. As an example of these observations, the following (arbitrarily-chosen) optimization problem is also a semidefinite program:

$$\begin{aligned} \underset{x,M,Y}{\text{minimize}}\ & a^\top x + \langle A_1, M\rangle + \langle A_2, Y\rangle \\ \text{subject to}\ & M + Y \succeq \Sigma \\ & \operatorname{diag}(M) = 1 \\ & \operatorname{tr}(Y) \leq 1 \\ & Y \succeq 0 \\ & \begin{bmatrix} 1 & x^\top \\ x & M \end{bmatrix} \succeq 0 \end{aligned} \tag{125}$$

(As a brief aside, the constraint $[1\ x^\top; x\ M] \succeq 0$ is equivalent to $xx^\top \preceq M$ which is in turn equivalent to $x^\top M^{-1} x \leq 1$ and $M \succeq 0$.)

Semidefinite constraints as quadratic polynomials. An alternative way of viewing the constraint $M \succeq 0$ is that the polynomial $p_M(v) = v^\top M v$ is non-negative for all $v \in \mathbb{R}^d$. More generally, if we have a non-hogoneous polynomial $p_{M,y,c}(v) = v^\top M v + y^\top v + c$, we have $p_{M,y,c}(v) \geq 0$ for all v if and only if $M' \succeq 0$ for $M' = \begin{bmatrix} c & y^\top/2 \\ y/2 & M \end{bmatrix} \succeq 0$.

This polynomial perspective is helpful for solving eigenvalue-type problems. For instance, $\|M\| \leq \lambda$ if and only if $v^\top M v \leq \lambda \|v\|_2^2$ for all v, which is eqvuialent to asking that $v^\top(\lambda I - M)v \geq 0$ for all v. Thus $\|M\|$ can be expressed as the solution to

$$\begin{aligned} &\text{minimize } \lambda \\ &\text{subject to } \lambda I - M \succeq 0 \text{ (equivalently, } v^\top(\lambda I - M)v \geq 0 \text{ for all } v) \end{aligned} \tag{126}$$

We thus begin to see a relationship between moments and *polynomial non-negativity constraints*.

Higher-degree polynomials. It is tempting to generalize the polynomial approach to higher moments. For instance, $M_4(p)$ denote the 4th moment tensor of p, i.e. the unique symmetric tensor such that

$$\langle M_4, v^{\otimes 4}\rangle = \mathbb{E}_{x\sim p}[\langle x - \mu, v\rangle^4]. \tag{127}$$

Note we can equivalently express $\langle M_4, v^{\otimes 4}\rangle = \sum_{ijkl} (M_4)_{ijkl} v_i v_j v_k v_l$, and hence $(M_4)_{ijkl} = \mathbb{E}[(x_i - \mu)(x_j - \mu)(x_k - \mu)(x_l - \mu)]$.

A distribution p has bounded 4th moment if and only if $\langle M, v^{\otimes 4}\rangle \leq \lambda \|v\|_2^4$ for all v. Letting $p_M(v) \stackrel{\text{def}}{=} \langle M, v^{\otimes 4}\rangle$, we thus can express the 4th moment of p as the polynomial program

$$\begin{aligned} &\text{minimize } \lambda \\ &\text{subject to } \lambda(v_1^2 + \cdots + v_d^2)^2 - p_M(v) \geq 0 \text{ for all } v \in \mathbb{R}^d \end{aligned} \tag{128}$$

Unfortunately, in constrast to (125), (128) is NP-hard to solve in general. We will next see a way to approximate (128) via a technique called *sum-of-squares programming*, which is a way of approximately reducing polynomial programs such as (128) to a large but polynomial-size semidefinite program.

Warm-up: certifying non-negativity over $\mathbb{R}$. Consider the one-dimensional polynomial

$$q(x) = 2x^4 + 2x^3 - x^2 + 5 \tag{129}$$

Is it the case that $q(x) \geq 0$ for all x? If so, how would we check this?

What if I told you that we had

$$q(x) = \frac{1}{2}(2x^2 + x - 3)^2 + \frac{1}{2}(3x + 1)^2 \tag{130}$$

Then, it is immediate that $q(x) \geq 0$ for all x, since it is a (weighted) sum of squares.

How can we construct such decompositions of q? First observe that we can re-write q as the matrix function

$$q(x) = \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix}^\top \underbrace{\begin{bmatrix} 5 & 0 & 0 \\ 0 & -1 & 1 \\ 0 & 1 & 2 \end{bmatrix}}_{M} \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix}. \tag{131}$$

On the other hand, the sum-of-squares decomposition for q implies that we can also write

$$q(x) = \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix}^\top \left(\frac{1}{2} \begin{bmatrix} -3 \\ 1 \\ 2 \end{bmatrix} \begin{bmatrix} -3 \\ 1 \\ 2 \end{bmatrix}^\top + \frac{1}{2} \begin{bmatrix} 1 \\ 3 \\ 0 \end{bmatrix} \begin{bmatrix} 1 \\ 3 \\ 0 \end{bmatrix}^\top \right) \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix}, \tag{132}$$

i.e. we can decompose the matrix M defining $q(x) = [1; x; x^2]^\top M [1; x; ^2]$ into a non-negative combination of rank-one outer products, which is true if and only if $M \succeq 0$.

There is one problem with this, which is that despite our successful decomposition of q, M is self-evidently not positive semidefinite! (For instance, $M_{22} = -1$.) The issue is that the matrix M defining $q(x)$ is not unique. Indeed, any $M(a) = \begin{bmatrix} 5 & 0 & -a \\ 0 & 2a-1 & 1 \\ -a & 1 & 2 \end{bmatrix}$ would give rise to the same $q(x)$, and a sum-of-squares decomposition merely implies that $M(a) \succeq 0$ for *some* a. Thus, we obtain the following characterization:

$$q(x) \text{ is a sum of squares } \sum_{j=1}^{k} q_j(x)^2 \iff M(a) \succeq 0 \text{ for some } a \in \mathbb{R}. \tag{133}$$

For the particular decomposition above we took $a = 3$.

Sum-of-squares in two dimensions. We can generalize the insights to higher-dimensional problems. Suppose for instance that we wish to check whether $q(x, y) = a_{40}x^4 + a_{31}x^3y + a_{22}x^2y^2 + a_{13}xy^3 + a_{04}y^4 + a_{30}x^3 + a_{21}x^2y + a_{12}xy^2 + a_{03}y^3 + a_{20}x^2 + a_{11}xy + a_{02}y^2 + a_{10}x + a_{01}y + a_{00}$ is non-negative for all x, y. Again, this is hard-to-check, but we can hope to check the sufficient condition that q is a sum-of-squares, which we will express as $q \succeq_{\text{sos}} 0$. As before this is equivalent to checking that a certain matrix is positive semidefinite. Observe that

$$q(x,y) = \begin{bmatrix} x^2 \\ xy \\ y^2 \\ x \\ y \\ 1 \end{bmatrix}^\top \begin{bmatrix} a_{40} & a_{31}/2 & -b & a_{30}/2 & -c & -b' \\ a_{31}/2 & a_{22}+2b & a_{13}/2 & a_{21}/2+c & -c' & -c'' \\ -b & a_{13}/2 & a_{04} & a_{21}/2+c' & a_{03}/2 & -b'' \\ a_{30}/2 & a_{21}/2+c & a_{21}/2+c' & a_{20}+2b' & a_{11}/2+c'' & a_{10}/2 \\ -c & -c' & a_{03}/2 & a_{11}/2+c'' & a_{02}+2b'' & a_{01}/2 \\ -b' & -c'' & -b'' & a_{10}/2 & a_{01}/2 & a_{00} \end{bmatrix} \begin{bmatrix} x^2 \\ xy \\ y^2 \\ x \\ y \\ 1 \end{bmatrix} \tag{134}$$

for any b, b', b'', c, c', c''. Call the above expression $M(b, b', b'', c, c', c'')$, which is linear in each of its variables. Then we have $q \succeq_{\text{sos}} 0$ if and only if $M(b, b', b'', c, c', c'') \succeq 0$ for some setting of the bs and cs.

Sum-of-squares in arbitrary dimensions. In general, if we have a polynomial $q(x_1, \ldots, x_d)$ in d variables, which has degree $2t$, then we can embed it as some matrix $M(b)$ (for decision variables b that capture the symmetries in M as above), and the dimensionality of M will be the number of monomials of degree at most t which turns out to be $\binom{d+t}{t} = \mathcal{O}((d+t)^t)$.

The upshot is that any constraint of the form $q \succeq_{\text{sos}} 0$, where q is linear in the decision variables, is a semidefinite constraint in disguise. Thus, we can solve any program of the form

$$\begin{aligned} &\underset{y}{\text{maximize}}\ c^\top y \\ &\text{subject to } q_1 \succeq_{\text{sos}} 0, \ldots, q_m \succeq_{\text{sos}} 0, \end{aligned} \tag{135}$$

where the q_j are linear in the decision variables y. (And we are free to throw in any additional linear inequality or semidefinite constraints as well.) We refer to such optimization problems as *sum-of-squares programs*, in analogy to semidefinite programs.

Sum-of-squares for kth moment. Return again to the kth moment problem. As a polynomial program we sought to minimize λ such that $\lambda(v_1^2 + \cdots + v_d^2)^{k/2} - \langle M_{2k}, v^{\otimes 2k}\rangle$ was a non-negative polynomial. It is then natural to replace the non-negativity constraint with the constraint that $\lambda\|v\|_2^k - \langle M_{2k}, v^{\otimes 2k}\rangle \succeq_{\text{sos}} 0$. However, we actually have a bit more flexibility and it turns out that the best program to use is

$$\begin{aligned} &\text{minimize } \lambda \\ &\text{subject to } \lambda - \langle M_{2k}, v^{\otimes 2k}\rangle + (\|v\|_2^2 - 1)q(v) \succeq_{\text{sos}} 0 \text{ for some } q \text{ of degree at most } 2k-2 \end{aligned} \tag{136}$$

Note that the family of all such q can be linearly parameterized and so the above is indeed a sum-of-squares program. It is always at least as good as the previous program because we can take $q(v) = \lambda(1 + \|v\|_2^2 + \cdots + \|v\|_2^{2k-2})$.

When the solution λ^* to (136) is at most σ^{2k} for $M_{2k}(p)$, we say that p has $2k$th moment *certifiably bounded* by σ^{2k}. In this case a variation on the filtering algorithm achieves error $\mathcal{O}(\sigma\epsilon^{1-1/2k})$. We will not discuss this in detail, but the main issue we need to resolve to obtain a filtering algorithm is to find some appropriate tensor T such that $\langle T, M_{2k}\rangle = \lambda^*$ and T "looks like" the expectation of $v^{\otimes 2k}$ for some probability distribution over the unit sphere. Then we can filter using $\tau_i = \langle T, (x_i - \hat{\mu})^{\otimes 2k}\rangle$.

To obtain T requires computing the dual of (136), which requires more optimization theory than we have assumed from the reader, but it can be done in polynomial time. We refer to the corresponding T as a *pseudomoment* matrix. Speaking very roughly, T has all properties of a moment matrix that can be "proved using only sum-of-squares inequalities", which includes all properties that we needed for the filtering algorithm to work. We will henceforth ignore the issue of T and focus on assumptions on p that ensure that $M_{2k}(p)$ is certifiably bounded. The main such assumption is the *Poincaré inequality*, which we cover in the next section.

[Lecture 11]

2.9 Sum-of-Squares Certifiably from the Poincaré inequality

We now turn our attention to bounding the value of (136). Ignoring finite-sample issues, our goal is to identify assumptions on p such that $M_{2k}(p) \stackrel{\text{def}}{=} \mathbb{E}_{X\sim p}[(X-\mu)^{\otimes 2k}]$ yields a small value for (136).

Before doing so, we will introduce some machinery for establishing bounds on (136). The main idea is that of a sum-of-squares proof:

Definition 2.46. A polynomial inequality $p(v) \le q(v)$ has a *sum-of-squares proof* if $q(v) - p(v) \succeq_{\text{sos}} 0$. We will also denote this as $q(v) \succeq_{\text{sos}} p(v)$ or $p(v) \preceq_{\text{sos}} q(v)$.

The usefulness of this perspective is that the relation $\preceq_{\text{sos}}$ satisfies many of the same properties as $\le$:

- If $p_1 \preceq_{\text{sos}} p_2$ and $p_2 \preceq_{\text{sos}} p_3$, then $p_1 \preceq_{\text{sos}} p_3$.
- If $p_1 \preceq_{\text{sos}} q_1$ and $p_2 \preceq_{\text{sos}} q_2$, then $p_1 + p_2 \preceq_{\text{sos}} q_1 + q_2$.
- If $p_1 \succeq_{\text{sos}} 0$ and $p_2 \succeq_{\text{sos}} 0$, then $p_1 p_2 \succeq_{\text{sos}} 0$.
- If $p_1 \preceq_{\text{sos}} p_2$, $q_1 \preceq_{\text{sos}} q_2$, and $p_2, q_1 \succeq_{\text{sos}} 0$, then $p_1 q_1 \preceq_{\text{sos}} p_2 q_2$.

- Moreover, many "standard" inequalities such as Cauchy-Schwarz and Hölder have sum-of-squares proofs.

Using these, we can often turn a normal proof that $p \leq q$ into a sum-of-squares proof that $p \preceq q$ as long as we give sum-of-squares proofs for a small number of key steps.

For concreteness, we will prove the last two claims properties above. We first prove that $p_1, p_2 \succeq_{\text{sos}} 0 \implies p_1 p_2 \succeq_{\text{sos}} 0$. Indeed we have

$$p_1(v)p_2(v) = (\sum_i p_{1i}(v)^2)(\sum_j p_{2j}(v)^2) = \sum_{ij}(p_1 i(v) p_{2j}(v))^2 \succeq_{\text{sos}} 0 \tag{137}$$

Next we prove that $p_1 \preceq_{\text{sos}} p_2, q_1 \preceq_{\text{sos}} q_2$, and $p_2, q_1 \succeq_{\text{sos}} 0$ implies $p_1 q_2 \preceq_{\text{sos}} p_2 q_2$. This is because

$$p_2 q_2 - p_1 q_1 = p_2(q_2 - q_1) + (p_2 - p_1)q_1 \succeq_{\text{sos}} 0, \tag{138}$$

where the second relation uses $p_2, q_2 - q_1 \succeq_{\text{sos}} 0$ and $p_2 - p_1, q_1 \succeq_{\text{sos}} 0$ together with the previous result.

In view of this, we can reframe bounding (136) as the following goal:

Goal: Find a sum-of-squares proof that $\langle M_{2k}(p), v^{\otimes 2k} \rangle \preceq_{\text{sos}} \lambda \|v\|_2^{2k}$.

Certifiability for Gaussians. We now return to the assumptions needed on p that will enable us to provide the desired sum-of-squares proof. Let us start by observing that a sum-of-squares proof exists for any Gaussian distribution: If $p = \mathcal{N}(\mu, \Sigma)$, then

$$\langle M_{2k}(\mathcal{N}(\mu, \Sigma)), v^{\otimes 2k} \rangle = \langle M_{2k}(\mathcal{N}(0, I)), (\Sigma^{1/2} v)^{\otimes 2k} \rangle \tag{139}$$

$$= \big(\prod_{i=1}^{k}(2i-1)\big)\langle \mathcal{I}, (\Sigma^{1/2} v)^{\otimes 2k} \rangle \tag{140}$$

$$= \big(\prod_{i=1}^{k}(2i-1)\big)\|\Sigma^{1/2} v\|_2^{2k} \tag{141}$$

$$\leq (2k)^k \|\Sigma\|^k \|v\|_2^{2k}, \tag{142}$$

so we may take $\lambda = (2k\|\Sigma\|)^k$. (Here $\mathcal{I}$ denotes the identity tensor that is 1 along the diagonal and zero elsewhere.) Therefore normal distributions have certifiably bounded moments, but the proof above heavily exploited the rotational symmetry of a normal distribution. We can provide similar proofs for other highly symmetric distributions (such as the uniform distribution on the hypercube), but these are unsatisfying as they only apply under very specific distributional assumptions. We would like more general properties that yield certifiably bounded moments.

Poincaré inequality. The property we will use is the *Poincaré inequality*. A distribution p on $\mathbb{R}^d$ is said to satisfy the Poincaré inequality with parameter σ if

$$\mathsf{Var}_{x \sim p}[f(x)] \leq \sigma^2 \mathbb{E}_{x \sim p}[\|\nabla f(x)\|_2^2] \tag{143}$$

for all differentiable functions $f : \mathbb{R}^d \to \mathbb{R}$. This is a "global to local property"–it says that for any function that for any function f that varies under p, that variation can be picked up in terms of local variation (the gradient). In particular, it says that p doesn't have any "holes" (regines with low probability density that lie between two regions of high probability density). Indeed, suppose that A and B were two disjoint convex regions with $p(A) = p(B) = \frac{1}{2}$. Then p cannot satisfy the Poincaré inequality with any constant, since there is a function that is 1 on A, 0 on B, and constant on both A and B.

Below are some additional examples and properties of Poincaré distributions:

- A one-dimensional Gaussian $\mathcal{N}(\mu, \sigma^2)$ is Poincaré with constant σ.
- If p, q are σ-Poincaré then their product $p \times q$ is σ-Poincaré. In particular a multivariate Gausssian $\mathcal{N}(\mu, \sigma^2 I)$ is σ-Poincaré.

- If $X \sim p$ is σ-Poincaré and A is a linear map, then AX is $(\sigma\|A\|)$-Poincaré. In particular, $aX_1 + aX_2$ is $(\sqrt{a^2+b^2}\sigma)$-Poincaré when X_1 and X_2 are both σ-Poincaré, and $\mathcal{N}(\mu, \Sigma)$ is $\|\Sigma\|^{1/2}$-Poincaré.
- More generally, if $X \sim p$ is σ-Poincaré and f is L-Lipschitz, then $f(X)$ is (σL)-Poincaré.

Together these imply that Poincaré distributions contain multivariate Gaussians, arbitrary Lipschitz functions of Gaussians, and independent sums of such distributions. The above properties (except the initial Gaussian property) are all straightforward computations. Let us next state two substantially deeper results:

- If p is σ-strongly log-concave (meaning that the log-probability density $\log p(x)$ satisfies $\nabla^2 \log p(x) \preceq -\frac{1}{\sigma^2} I$), then p is σ-Poincaré (Bakry and Émery, 1985).
- Suppose that the support of $X \sim p$ has ℓ_2-radius at most R, and let $Z = \mathcal{N}(0, \tau^2 I)$ for $\tau \geq 2R$. Then $X + Z$ is $(\tau\sqrt{e})$-Poincaré (Bardet et al., 2018).

Thus Poincaré encompasses all strongly log-concave densities, and effectively any product of bounded random variables (after adding Gaussian noise, which we can always do ourselves).

It is instructive to compare Poincaré to the sub-Gaussian property that we have so far relied on. Poincaré is neither strictly stronger or weaker than sub-Gaussian, but it is stronger than sub-exponential (we will see this below). In general, we should think of Poincaré as being substantially stronger than sub-exponential: it implies that not only is the distribution itself sub-exponential, but so is any Lipschitz function of the density.

As an example, consider the random variable $(X, Y) \in \mathbb{R}^d$ where $X \sim \mathcal{N}(0, I)$ and $Y = \epsilon X$ for a Rademacher random variable ϵ. Then (X, Y) is sub-Gaussian, but not Poincaré with good constant: if we take $f(X, Y) = \sum_i X_i Y_i$, then f is with high probability close to either $+d$ or $-d$, so $\mathsf{Var}[f(X, Y)] \approx d^2$. However, $\nabla f(X, Y) = (Y_1, \ldots, Y_d, X_1, \ldots, X_d)$ and so $\|\nabla f(X, Y)\|_2^2$ is close to $2d$ with high probability. Thus while the sub-Gaussian constant is $\mathcal{O}(1)$, the Poincaré constant in this case is $\Omega(\sqrt{d})$.

Consequences of Poincaré. So far we have seen conditions that imply Poincaré, but we would also like to derive consequences of this property. Below are some of the most useful ones:

- If $X \sim p$ is σ-Poincaré, then Lipschitz functions concentrate: $\mathbb{P}[|f(x) - \mathbb{E}[f(x)]| \geq t] \leq 6\exp(-t/(\sigma L))$ for any L-Lipschitz f.
- As a corollary, we have *volume expansion*: For any set A, let A_ϵ be the set of points within ℓ_2-distance ϵ of A. Then $p(A)p(A_\epsilon^c) \leq 36\exp(-\epsilon/\sigma)$.

This second property implies, for instance, that if $p(A) \geq \delta$, then almost all points will be within distance $\mathcal{O}(\sigma \log(1/\delta))$ of A.

To prove the second property, let $f(x) = \min(\inf_{y \in A} \|x - y\|_2 . \epsilon)$. Then f is Lipschitz, is 0 on A, and is ϵ on A_ϵ^c. Let μ be the mean of $f(X)$. Since f is sub-exponential we have $p(A) = p(f(X) = 0) \leq 6\exp(-\mu/\sigma)$, and $p(A_\epsilon^c) = p(f(X) = \epsilon) \leq 6\exp(-(\epsilon - \mu))/\sigma)$. Multiplying these together yields the claimed result.

The most important property for our purposes, however, will be the following:

Theorem 2.47. *Suppose that p is σ-Poincaré and let f be a differentiable function such that $\mathbb{E}[\nabla^j f(X)] = 0$ for $j = 1, \ldots, k-1$. Then there is a universal constant C_k such that $\mathsf{Var}[f(X)] \leq C_k \sigma^{2k} \mathbb{E}[\|\nabla^k f(X)\|_F^2]$.*

Note that $k = 1$ is the original Poincaré property, so we can think of Theorem 2.47 as a generalization of Poincaré to higher derivatives. Note also that $\nabla^k f(X)$ is a tensor in $\mathbb{R}^{d^k}$; the notation $\|\nabla^k f(X)\|_F^2$ denotes the squared Frobenius norm of $\nabla^k f(X)$, i.e. the sum of the squares of its entries.

Theorem 2.47, while it may appear to be a simple generalization of the Poincaré property, is a deep result that was established in Adamczak and Wolff (2015), building on work of Latała (2006). We will use Theorem 2.47 in the sequel to construct our sum-of-squares proofs.

Sum-of-squares proofs for Poincaré distributions. Here we will construct sum-of-squares proofs that $M_{2k}(v) \stackrel{\text{def}}{=} \mathbb{E}_p[\langle x - \mu, v\rangle^{2k}] \preceq_{\text{sos}} C_k' \sigma^{2k} \|v\|_2^{2k}$ whenever p is σ-Poincaré, for some universal constants C_k'. We

will exhibit the proof for $k = 1, 2, 3$ (the proof extends to larger k and the key ideas appear already by $k = 3$). We introduce the notation

$$M_k = \mathbb{E}[(x-\mu)^{\otimes k}], \tag{144}$$

$$M_k(v) = \langle M_k, v^{\otimes k}\rangle = \mathbb{E}[\langle x-\mu, v\rangle^k]. \tag{145}$$

Proof for $k = 1$. We wish to show that $\mathbb{E}_p[\langle x-\mu, v\rangle^2] \preceq_{\text{sos}} \sigma^2\|v\|_2^2$. To do this take $f_v(x) = \langle x, v\rangle$. Then the Poincaré inequality applied to f_v yields

$$\mathbb{E}_p[\langle x-\mu, v\rangle^2] = \mathsf{Var}[f_v(x)] \le \sigma^2\mathbb{E}[\|\nabla f_v(x)\|_2^2] = \sigma^2\mathbb{E}[\|v\|_2^2] = \sigma^2\|v\|_2^2. \tag{146}$$

Thus $M_2(v) \le \sigma^2\|v\|_2^2$ (this is just saying that Poincaré distributions have bounded covariance). This property has a sum-of-squares proof because it is equivalent to $\sigma^2 I - M_2 \succeq 0$, and we know that all positive semidefiniteness relations are sum-of-squares certifiable.

Proof for $k = 2$. Extending to $k = 2$, it makes sense to try $f_v(x) = \langle x-\mu, v\rangle^2$. Then we have $\nabla f_v(x) = 2\langle x-\mu, v\rangle v$ and hence $\mathbb{E}[\nabla f_v(x)] = 0$. We also have $\nabla^2 f_v(x) = 2v\otimes v$. Thus applying Theorem 2.47 we obtain

$$\mathsf{Var}[f_v(x)] \le C_2\sigma^4\mathbb{E}[\|2v\otimes v\|_F^2] = 4C_2\sigma^4\|v\|_2^4. \tag{147}$$

We also have $\mathsf{Var}[f_v(x)] = \mathbb{E}[\langle x-\mu, v\rangle^4] - \mathbb{E}[\langle x-\mu, v\rangle^2]^2 = M_4(v) - M_2(v)^2$. Thus

$$M_4(v) = (M_4(v) - M_2(v)^2) + M_2(v)^2 \tag{148}$$

$$\le 4C_2\sigma^4\|v\|_2^4 + \sigma^4\|v\|_2^4 = (4C_2+1)\sigma^4\|v\|_2^4. \tag{149}$$

This shows that the fourth moment is bounded, but how can we construct a sum-of-squares proof? We already have that $M_2(v)^2 \preceq_{\text{sos}} \sigma^4\|v\|_2^4$ (by $0 \preceq_{\text{sos}} M_2(v) \preceq_{\text{sos}} \sigma^2\|v\|_2^2$ and the product property). Therefore we focus on bounding $M_4(v) - M_2(v)^2 = \mathsf{Var}[f_v(x)]$.

For this we will apply Theorem 2.47 to a modified version of $f_v(x)$. For a matrix A, let $f_A(x) = (x-\mu)^\top A(x-\mu) = \langle A, (x-\mu)^{\otimes 2}\rangle$. Then $f_v(x) = f_A(x)$ for $A = vv^\top$. By the same calculations as above we have $\mathbb{E}[\nabla f_A(x)] = 0$ and $\nabla^2 f_A(x) = 2A$. Thus by Theorem 2.47 we have

$$\mathsf{Var}[f_A(x)] \le C_2\sigma^4\mathbb{E}[\|2A\|_F^2] = 4C_2\sigma^4\|A\|_F^2. \tag{150}$$

On the other hand, we have $\mathsf{Var}[f_A(x)] = \langle M_4, A\otimes A\rangle - \langle M_2, A\rangle^2 = \langle M_4 - M_2\otimes M_2, A\otimes A\rangle$. Thus (150) implies that

$$\langle M_4 - M_2\otimes M_2, A\otimes A\rangle \le 4C_2\sigma^4\|A\|_F^2. \tag{151}$$

Another way of putting this is that $M_4 - M_2\otimes M_2$, when considered as a matrix in $\mathbb{R}^{d^2\times d^2}$, is smaller than $4C_2\sigma^4 I$ in the semidefinite ordering. Hence $4C_2\sigma^4 I - (M_4 - M_2\otimes M_2) \succeq 0$ and so $4C_2\sigma^4\|v\|_2^4 - \langle M_4 - M_2\otimes M_2, v^{\otimes 4}\rangle \preceq_{\text{sos}} 0$, giving us our desired sum-of-squares proof. To recap, we have:

$$M_4(v) = (M_4(v) - M_2(v)^2) + M_2(v)^2 \tag{152}$$

$$\preceq_{\text{sos}} 4C_2\sigma^4\|v\|_2^4 + \sigma^4\|v\|_2^4 = (4C_2+1)\sigma^4\|v\|_2^4, \tag{153}$$

so we can take $C_2' = 4C_2 + 1$.

Proof for $k = 3$. Inspired by the $k = 1, 2$ cases, we try $f_v(x) = \langle x-\mu, v\rangle^3$. However, this choice runs into problems, because $\nabla f_v(x) = 3\langle x-\mu, v\rangle^2 v$ and so $\mathbb{E}[\nabla f_v(x)] = 3M_2(v)v \neq 0$. We instead should take

$$f_v(x) = \langle x-\mu, v\rangle^3 - 3M_2(v)\langle x-\mu, v\rangle, \text{ which yields} \tag{154}$$

$$\mathbb{E}[\nabla f_v(x)] = \mathbb{E}[3\langle x-\mu, v\rangle^2 v - 3M_2(v)v] = 0, \tag{155}$$

$$\mathbb{E}[\nabla^2 f_v(x)] = \mathbb{E}[6\langle x-\mu, v\rangle(v\otimes v)] = 0, \tag{156}$$

$$\nabla^3 f_v(x) = 6(v\otimes v\otimes v). \tag{157}$$

Applying Theorem 2.47 to $f_v(x)$ yields

$$\mathsf{Var}[f_v(x)] \le C_3\sigma^6\|6(v\otimes v\otimes v)\|_F^2 = 36C_3\sigma^6\|v\|_2^6. \tag{158}$$

We can additionally compute

$$\mathsf{Var}[f_v(x)] = \mathbb{E}[(\langle x-\mu, v\rangle^3 - 3M_2(v)\langle x-\mu, v\rangle)^2] - \mathbb{E}[\langle x-\mu, v\rangle^3 - 3M_2(v)\langle x-\mu, v\rangle]^2 \tag{159}$$

$$= M_6(v) - 6M_2(v)M_4(v) + 9M_2(v)^3 - M_3(v)^2. \tag{160}$$

Since our goal is to bound $M_6(v)$, we re-arrange to obtain

$$M_6(v) = \mathsf{Var}[f_v(x)] + 6M_2(v)M_4(v) + M_3(v)^2 - 9M_2(v)^2 \tag{161}$$

$$\le 36C_3\sigma^6\|v\|_2^6 + 6(\sigma^2\|v\|_2^2)(C_2'\sigma^4\|v\|_2^4) + M_3(v)^2 + 0 \tag{162}$$

We can also use Hölder's inequality to obtain $M_3(v)^2 \le M_2(v)M_4(v)$, which yields an overall bound of $M_6(v) \le (36C_3 + 12C_2')\sigma^6\|v\|_2^6$.

We now turn this into a sum-of-squares proof. We need to show the following four relations:

$$(i)\ \mathsf{Var}[f_v(x)] \preceq_{\text{sos}} 36C_3\sigma^6\|v\|_2^6, \quad (ii)\ M_2(v)M_4(v) \preceq_{\text{sos}} (\sigma^2\|v\|_2^2)(C_2'\sigma^4\|v\|_2^4), \tag{163}$$

$$(iii)\ M_3(v) \preceq_{\text{sos}} M_2(v)M_4(v), \quad (iv)\ -9M_2(v)^2 \preceq_{\text{sos}} 0. \tag{164}$$

The relation (ii) again follows by the product property of $\preceq_{\text{sos}}$, while $-9M_2(v)^2 \preceq_{\text{sos}} 0$ is direct because $M_2(v)^2$ is already a square. We will show in an exercise that the Hölder inequality in (iii) has a sum-of-squares proof, and focus on (i).

The relation (i) holds for reasons analogous to the $k = 2$ case. For a symmetric tensor $A \in \mathbb{R}^{d^3}$, let $f_A(x) = \langle A, (x-\mu)^{\otimes 3} - 3M_2 \otimes (x-\mu)\rangle$. Then just as before we have $\mathbb{E}[\nabla f_A(x)] = 0$, $\mathbb{E}[\nabla^2 f_A(x)] = 0$, and so $\mathsf{Var}[f_A(x)] \le 36C_3\sigma^6\|A\|_F^2$, which implies that[3]

$$M_6 - 6M_2 \otimes M_4 + 9M_2 \otimes M_2 \otimes M_2 - M_3 \otimes M_3 \preceq 36C_3\sigma^6 I, \tag{165}$$

and hence $\mathsf{Var}[f_v(x)] \preceq_{\text{sos}} 36C_3\sigma^6\|v\|_2^6$ (again because semidefinite relations have sum-of-squares proofs).

In summary, we have $M_6(v) \preceq_{\text{sos}} (36C_3 + 12C_2')\sigma^6\|v\|_2^6$, as desired.

Generalizing to higher k. For higher k the proof is essentially the same. What is needed is a function $f_v(x)$ whose first $k-1$ derivates all have zero mean. This always exists and is unique up to scaling by constants. For instance, when $k = 4$ we can take $f_v(x) = \langle x-\mu, v\rangle^4 - 6M_2(v)\langle x-\mu, v\rangle^2 - 4M_3(v)\langle x-\mu, v\rangle - M_4(v) + 6M_2(v)^2$. This appears somewhat clunky but is a special case of a combinatorial sum. For the general case, let $\mathcal{T}_k$ be the set of all integer tuples $(i_0, i_1, \ldots)$ such that $i_0 \ge 0$, $i_s \ge 2$ for $s > 0$, and $i_0 + i_1 + \cdots = k$. Then the general form is

$$f_{v,k}(x) = \sum_{(i_0,\ldots,i_r)\in\mathcal{T}_k} (-1)^r \binom{k}{i_0\ \cdots\ i_r} \langle x-\mu, v\rangle^{i_0} M_{i_1}(v)M_{i_2}(v)\cdots M_{i_r}(v). \tag{166}$$

The motivation for this formula is that it is the solution to $\nabla f_{v,k}(x) = k f_{v,k-1}(x)v$. Using $f_{v,k}$, one can construct sum-of-squares proofs by applying Theorem 2.47 to the analogous $f_{A,k}$ function as before, and then use induction, the product rule, and Hölder's inequality as in the $k = 3$ case.

[Lectures 12-13]

3 Resilience Beyond Mean Estimation

We have so far focused primarily on mean estimation, first considering information theoretic and then algorithmic issues. We now turn back to information theoretic issues with a focus on generalizing our results from mean estimation to other statistical problems.

Let us recall our general setup: for true (test) distribution p^* and corrupted (train) distribution $\tilde{p}$, we observe samples $X_1, \ldots, X_n$ from $\tilde{p}$ (oblivious contamination, although we can also consider adaptive

[3] Actually this is not quite true because we only bound $\mathsf{Var}[f_A(x)]$ for symmetric tensors A. What is true is that this holds if we symmetrize the left-hand-side of (165), which involves averaging over all ways of splitting M_2 and M_4 over the 3 copies of $\mathbb{R}^d$ in $\mathbb{R}^{d\times d\times d}$.

contamination as in Section 2.6.2). We wish to estimate a parameter θ and do so via en estimator $\hat{\theta} = \hat{\theta}(X_1, \ldots, X_n)$. Our goal is to construct an estimator such that $L(p^*, \hat{\theta})$ is small according to a given loss function L. This was summarized in Figure 1 from Section 1.

As before, we will start by allowing our estimator $\hat{\theta}$ to directly access the population distribution $\tilde{p}$ rather than samples. Thus we wish to control the error $L(p^*, \hat{\theta}(\tilde{p}))$. Since this is hopeless without further assumptions, we assume that $D(p^*, \tilde{p}) \le \epsilon$ for some distance D, and that p^* lies in some family $\mathcal{G}$.

For now we continue to take $D = \mathsf{TV}$ and focus on more general losses L, corresponding to tasks beyond mean estimation. Two key examples will be:

- **Second moment estimation** in spectral norm, which corresponds to the loss $L(p, S) = \|\mathbb{E}_p[XX^\top] - S\|$.
- **Linear regression**, which corresponds to the loss $L(p, \theta) = \mathbb{E}_{x,y\sim p}[(y - \theta^\top x)^2 - (y - \theta^*(p)^\top x)^2]$. Note that here L measures the *excess predictive loss* so that $L(p, \theta^*(p)) = 0$.

As in the mean estimation case, we will define the modulus of continuity and the family of resilience distributions, and derive sufficient conditions for resilience.

Modulus of continuity. The modulus of continuity generalizes straightforwardly from the mean estimation case. We define

$$\mathfrak{m}(\mathcal{G}, 2\epsilon, L) = \sup_{p,q\in\mathcal{G}:\mathsf{TV}(p,q)\le 2\epsilon} L(p, \theta^*(q)). \tag{167}$$

As before, the modulus $\mathfrak{m}$ upper-bounds the minimax loss. Specifically, consider the projection estimator that outputs $\hat{\theta}(\tilde{p}) = \theta^*(q)$ for any $q \in \mathcal{G}$ with $\mathsf{TV}(\tilde{p}, q) \le \epsilon$. Then the error of $\hat{\theta}$ is at most $\mathfrak{m}$ because $\mathsf{TV}(q, p^*) \le 2\epsilon$ and $p^*, q \in \mathcal{G}$.

Resilience. Generalizing resilience requires more care. Recall that for mean estimation the set of (ρ, ϵ)-resilient distributions was

$$\mathcal{G}^{\mathsf{TV}}_{\mathsf{mean}}(\rho, \epsilon) \stackrel{\text{def}}{=} \big\{p \mid \|\mathbb{E}_r[X] - \mathbb{E}_p[X]\| \le \rho \text{ for all } r \le \frac{p}{1-\epsilon}\big\}. \tag{168}$$

We saw in Section 2.4 that robust mean estimation is possible for the family $\mathcal{G}_{\mathsf{mean}}$ of resilient distributions; the two key ingredients were the existence of a midpoint distribution and the triangle inequality for $L(p, \theta^*(q)) = \|\mu_p - \mu_q\|$. We now extend the definition of resilience to arbitrary cost functions $L(p, \theta)$ that may not satisfy the triangle inequality. The general definition below imposes two conditions: (1) the parameter $\theta^*(p)$ should do well on all distributions $r \le \frac{p}{1-\epsilon}$, and (2) any parameter that does well on some $r \le \frac{p}{1-\epsilon}$ also does well on p. We measure performance on r with a *bridge function* $B(r, \theta)$, which is often the same as the loss L but need not be.

Definition 3.1 ($\mathcal{G}^{\mathsf{TV}}(\rho_1, \rho_2, \epsilon)$)**.** Given an arbitrary loss function $L(p, \theta)$, we define $\mathcal{G}^{\mathsf{TV}}(\rho_1, \rho_2, \epsilon) = \mathcal{G}^{\mathsf{TV}}_{\downarrow}(\rho_1, \epsilon) \cap \mathcal{G}^{\mathsf{TV}}_{\uparrow}(\rho_1, \rho_2, \epsilon)$, where:

$$\mathcal{G}^{\mathsf{TV}}_{\downarrow}(\rho_1, \epsilon) \triangleq \{p \mid \sup_{r\le\frac{p}{1-\epsilon}} B(r, \theta^*(p)) \le \rho_1\}, \tag{169}$$

$$\mathcal{G}^{\mathsf{TV}}_{\uparrow}(\rho_1, \rho_2, \epsilon) \triangleq \{p \mid \text{ for all } \theta, r \le \frac{p}{1-\epsilon}, (B(r, \theta) \le \rho_1 \Rightarrow L(p, \theta) \le \rho_2)\}, \tag{170}$$

The function $B(p, \theta)$ is an arbitrary cost function that serves the purpose of bridging.

If we take $B(p, \theta) = L(p, \theta) = \|\mathbb{E}_p[X] - \mathbb{E}_\theta[X]\|, \rho_2 = 2\rho_1$, then this exactly reduces to the resilient set $\mathcal{G}^{\mathsf{TV}}_{\mathsf{mean}}(\rho_1, \epsilon)$ for mean estimation. To see the reduction, note that $\mathcal{G}^{\mathsf{TV}}_{\mathsf{mean}}$ is equivalent to $\mathcal{G}^{\mathsf{TV}}_{\downarrow}$ in Equation (169). Thus we only need to show that $\mathcal{G}^{\mathsf{TV}}_{\uparrow}$ is a subset of $\mathcal{G}^{\mathsf{TV}}_{\downarrow}$. By our choice of B, L and ρ_2, the implication condition in $\mathcal{G}^{\mathsf{TV}}_{\uparrow}$ follows from the triangle inequality..

We will show that $\mathcal{G}^{\mathsf{TV}}$ is *not too big* by bounding its modulus of continuity, and that it is *not too small* by exhibiting reasonable sufficient conditions for resilience.

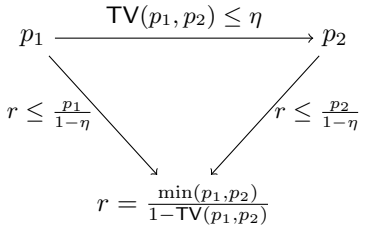

$$p_1 \in \mathcal{G}_{\downarrow}^{\mathsf{TV}} \implies B(r, \theta^*(p_1)) \le \rho_1 \xRightarrow{p_2 \in \mathcal{G}_{\uparrow}^{\mathsf{TV}}} L(p_2, \theta^*(p_1)) \le \rho_2$$

Figure 8: Midpoint distribution helps bridge the modulus for $\mathcal{G}^{\mathsf{TV}}$.

Not too big: bounding $\mathfrak{m}$. We show that the designed $\mathcal{G}^{\mathsf{TV}}(\rho_1, \rho_2, \epsilon)$ has small modulus of continuity (and thus population minimax limit) in the following theorem:

Theorem 3.2. *For $\mathcal{G}^{\mathsf{TV}}(\rho_1, \rho_2, \epsilon)$ in Definition 3.1, we have $\mathfrak{m}(\mathcal{G}^{\mathsf{TV}}(\rho_1, \rho_2, \epsilon), \epsilon) \le \rho_2$.*

Proof. As illustrated in Figure 8, we still rely on the midpoint distribution r to bridge the modulus. Consider any p_1, p_2 satisfying $\mathsf{TV}(p_1, p_2) \le \epsilon$. Then there is a midpoint r such that $r \le \frac{p_1}{1-\epsilon}$ and $r \le \frac{p_2}{1-\epsilon}$. From the fact that $p_1 \in \mathcal{G}^{\mathsf{TV}}(\rho_1, \rho_2, \epsilon) \subset \mathcal{G}_{\downarrow}^{\mathsf{TV}}(\rho_1, \epsilon)$, we have $B(r, \theta^*(p_1)) \le \rho_1$. From this and the fact that $p_2 \in \mathcal{G}^{\mathsf{TV}}(\rho_1, \rho_2, \epsilon) \subset \mathcal{G}_{\uparrow}^{\mathsf{TV}}(\rho_1, \rho_2, \epsilon)$, we then have $L(p_2, \theta^*(p_1)) \le \rho_2$. Since p_1 and p_2 are arbitrary, this bounds the modulus of continuity by ρ_2. □

Not too small: concrete examples. We next show that $\mathcal{G}^{\mathsf{TV}}$ yields sensible conditions for second moment estimation and linear regression. We start with second moment estimation:

Proposition 3.3. *Let $B(p, S) = L(p, S) = \|\mathbb{E}_p[XX^\top] - S\|$, and let p be a distribution on $\mathbb{R}^d$ such that $p \in \mathcal{G}_{\mathsf{mom},k})(\sigma)$, i.e. p^* has bounded kth moments. Then assuming $k > 2$, we have $p \in \mathcal{G}^{\mathsf{TV}}(\rho, 2\rho, \epsilon)$ for $\rho = \mathcal{O}(\sigma^2 \epsilon^{1-2/k})$.*

This is essentially the same statement as for mean estimation, except with $\sigma^2 \epsilon^{1-2/k}$ instead of $\sigma \epsilon^{1-1/k}$.

Proof. First we show that $p \in \mathcal{G}^{\downarrow}(\rho, \epsilon)$, for which we need to show that

$$\|\mathbb{E}_r[XX^\top] - \mathbb{E}_p[XX^\top]\| \le \rho \text{ for all } r \le \frac{p}{1-\epsilon}. \tag{171}$$

Letting $Y = XX^\top$, this asks that Y is resilient in operator norm, which in turn asks that $\langle Y, Z \rangle$ is resilient for any $\|Z\|_* \le 1$, where $\|\cdot\|_*$ is dual to the operator norm. Recalling that the operator norm is the maximum singular value, it turns out that $\|\cdot\|_*$ is the *nuclear norm*, or the sum of the singular values. Thus for $Z = U\Lambda V^\top$ we have $\|Z\|_* = \sum_i \Lambda_{ii}$. (Proving this duality requires some non-trivial but very useful matrix inequalities that we provide at the end of this section.)

Conveniently, the extreme points of the nuclear norm ball are exactly rank-one matrices of the form $\pm vv^\top$ where $\|v\|_2 = 1$. Thus we exactly need that $\langle v, X \rangle^2$ is resilience for all v. Fortunately we have that $\mathbb{E}[|\langle v, X\rangle^2 - \mathbb{E}[\langle v, X\rangle^2]|^{k/2}] \le \mathbb{E}[|\langle v, X\rangle|^k] \le \sigma^k$, so p is (ρ_1, ϵ)-resilient with $\rho_1 = \sigma^2 \epsilon^{1-2/k}$, which gives that $p \in \mathcal{G}^{\downarrow}$.

Next we need to show that $p \in \mathcal{G}^{\uparrow}$. We want

$$\|\mathbb{E}_r[XX^\top] - S\| \le \rho_1 \implies \|\mathbb{E}_p[XX^\top] - S\| \le \rho_2 \text{ whenever } r \le \frac{p}{1-\epsilon}, \tag{172}$$

but this is the same as $\rho_2 - \rho_1 \le \|\mathbb{E}_r[XX^\top] - \mathbb{E}_p[XX^\top]\|$, and we already know that the right-hand-side is bounded by ρ_1, so we can take $\rho_2 = 2\rho_1$, which proves the claimed result. □

We move on to linear regression. In the proof for second moment estimation, we saw that $p \in \mathcal{G}^{\uparrow}$ was essentially implied by $p \in \mathcal{G}^{\downarrow}$. This was due to the symmetry of the second moment loss together with the

triangle inequality for $\|\cdot\|$, two properties that we don't have in general. The proof for second moment estimation will require somewhat more different proofs for $\mathcal{G}^{\uparrow}$ and $\mathcal{G}^{\downarrow}$. For simplicity we state the result only for fourth moments:

Proposition 3.4. *For a distribution p on $\mathbb{R}^d \times \mathbb{R}$, let $B(p,\theta) = L(p,\theta) = \mathbb{E}_p[(y - \langle\theta, x\rangle)^2 - (y - \langle\theta^*(p), x\rangle)^2]$. Let $Z = Y - \langle\theta^*(p), X\rangle$ and suppose that the following two conditions holds:*

$$\mathbb{E}_p[XZ^2X^\top] \preceq \sigma^2\mathbb{E}[XX^\top], \tag{173}$$

$$\mathbb{E}_p[\langle X, v\rangle^4] \le \kappa\mathbb{E}_p[\langle X, v\rangle^2]^2 \text{ for all } v. \tag{174}$$

Then $p \in \mathcal{G}^{\mathsf{TV}}(\rho, 5\rho, \epsilon)$ for $\rho = 2\sigma^2\epsilon$ as long as $\epsilon(\kappa - 1) \le \frac{1}{6}$ and $\epsilon \le \frac{1}{8}$.

Let us interpret the two conditions. First, as long as X and Z are independent (covariates are independent of noise), we have $\mathbb{E}_p[XZ^2X^\top] = \mathbb{E}[Z^2]\mathbb{E}[XX^\top]$, so in that case σ^2 is exactly a bound on the noise Z. Even when X and Z are not independent, the first condition holds when Z has bounded 4th moment.

The second condition is a *hypercontractivity condition* stating that the fourth moments of X should not be too large compared to the second moments. It is a bit unusual from the perspective of mean estimation, because it does not require X to be well-concentrated, but only well-concentrated relative to its variance. For regression, this condition makes sense because κ bounds how close the covariates are to being rank-deficient (the worst-case is roughly an ϵ-mass at some arbitrary distance $t/\sqrt{\epsilon}$, which would have second moment t^2 and fourth moment t^4/ϵ, so we roughly want $\kappa < 1/\epsilon$). We will show later that such a hypercontractivity condition is needed, i.e. simply assuming sub-Gaussianity (without making it relative to the variance) allows for distributions that are hard to robustly estimate due to the rank-deficiency issue.

Proof. First note that $L(p,\theta) = (\theta - \theta^*(p))^\top S_p(\theta - \theta^*(p))$, where $S_p = \mathbb{E}_p[XX^\top]$, and analogously for $L(r,\theta)$. At a high level our strategy will be to show that $\theta^*(r) \approx \theta^*(p)$ and $S_r \approx S_p$, and then use this to establish membership in $\mathcal{G}^{\downarrow}$ and $\mathcal{G}^{\uparrow}$.

We first use the hypercontractivity condition to show that $S_r \approx S_p$. We have

$$\mathbb{E}_r[\langle v, X\rangle^2] \ge \mathbb{E}_p[\langle v, X\rangle^2] - \frac{1}{1-\epsilon}\sqrt{\epsilon\mathsf{Var}_p[\langle v, X\rangle^2]} \tag{175}$$

$$= \mathbb{E}_p[\langle v, X\rangle^2] - \frac{1}{1-\epsilon}\sqrt{\epsilon(\mathbb{E}_p[\langle v, X\rangle^4] - \mathbb{E}_p[\langle v, X\rangle^2]^2)} \tag{176}$$

$$\ge \mathbb{E}_p[\langle v, X\rangle^2] - \frac{1}{1-\epsilon}\sqrt{\epsilon(\kappa-1)}\mathbb{E}_p[\langle v, X\rangle^2] \tag{177}$$

$$= (1 - \frac{1}{1-\epsilon}\sqrt{\epsilon(\kappa-1)})\mathbb{E}_p[\langle v, X\rangle^2]. \tag{178}$$

Thus $S_r \succeq (1 - \frac{1}{1-\epsilon}\sqrt{\epsilon(\kappa-1)})S_p$, and similarly $S_r \preceq (1 + \frac{1}{1-\epsilon}\sqrt{\epsilon(\kappa-1)})S_p$. Assuming $\epsilon \le \frac{1}{8}$ and $\epsilon(\kappa-1) \le \frac{1}{6}$, we have $\frac{1}{1-\epsilon}\sqrt{\epsilon(\kappa-1)} \le \frac{8}{7}\sqrt{1/6} < \frac{1}{2}$, and so $\frac{1}{2}S_p \preceq S_r \preceq \frac{3}{2}S_p$.

We now turn to $\mathcal{G}^{\uparrow}$ and $\mathcal{G}^{\downarrow}$. A useful relation is $\theta^*(p) = S_p^{-1}\mathbb{E}_p[XY]$, and $\theta^*(r) - \theta^*(p) = S_r^{-1}\mathbb{E}_r[XZ]$. To prove that $p \in \mathcal{G}^{\downarrow}$ we need to show that $(\theta^*(r) - \theta^*(p))^\top S_r(\theta^*(r) - \theta^*(p))$ is small. We have

$$(\theta^*(r) - \theta^*(p))^\top S_r(\theta^*(r) - \theta^*(p)) \le \frac{3}{2}(\theta^*(r) - \theta^*(p))^\top S_p(\theta^*(r) - \theta^*(p)) \tag{179}$$

$$= \frac{3}{2}\mathbb{E}_r[XZ]^\top S_p^{-1}\mathbb{E}_r[XZ] = \frac{3}{2}\|\mathbb{E}_r[S_p^{-1/2}XZ] - \mathbb{E}_p[S_p^{-1/2}XZ]\|_2^2. \tag{180}$$

This final condition calls for $S_p^{-1/2}XZ$ to be resilient, and bounded variance of this distribution can be seen to exactly correspond to the condition $\mathbb{E}[XZ^2X^\top] \preceq \sigma^2\mathbb{E}[XX^\top]$. Thus we have resilience with $\rho = \frac{3\sigma^2\epsilon}{2(1-\epsilon)^2} \le 2\sigma^2\epsilon$ (since $\epsilon < \frac{1}{8}$).

Now we turn to $\mathcal{G}^{\uparrow}$. We want that $(\theta - \theta^*(r))^\top S_r(\theta - \theta^*(r)) \le \rho$ implies $(\theta - \theta^*(p))^\top S_p(\theta - \theta^*(p)) \le 5\rho$.

By the triangle inequality we have

$$\sqrt{(\theta-\theta^*(p))^\top S_p(\theta-\theta^*(p))} \le \sqrt{(\theta-\theta^*(r))^\top S_p(\theta-\theta^*(r))} + \sqrt{(\theta^*(r)-\theta^*(p))^\top S_p(\theta^*(r)-\theta^*(p))} \quad (181)$$

$$\le \sqrt{2(\theta-\theta^*(r))^\top S_r(\theta-\theta^*(r))} + \sqrt{(4/3)\sigma^2\epsilon} \quad (182)$$

$$\le \sqrt{2\rho} + \sqrt{(4/3)\sigma^2\epsilon} = \sqrt{\rho}(\sqrt{2}+\sqrt{2/3}) \le \sqrt{5\rho}, \quad (183)$$

which completes the proof. □

Lower bound. TBD

Proving that nuclear norm is dual to operator norm. Here we establish a series of matrix inequalities that are useful more broadly, and use these to analyze the nuclear norm. The first allows us to reduce dot products of arbitrary matrices to symmetric PSD matrices:

Proposition 3.5. *For any (rectangular) matrices A, B of equal dimensions, we have*

$$\langle A, B\rangle^2 \le \langle (A^\top A)^{1/2}, (B^\top B)^{1/2}\rangle \langle (AA^\top)^{1/2}, (BB^\top)^{1/2}\rangle. \quad (184)$$

In a sense, this is like a "matrix Cauchy-Schwarz".

Proof. We first observe that $\begin{bmatrix} (AA^\top)^{1/2} & A \\ A^\top & (A^\top A)^{1/2} \end{bmatrix} \succeq 0$. This is because, if $A = U\Lambda V^\top$ is the singular value decomposition, we can write the above matrix as $\begin{bmatrix} U\Lambda U^\top & U\Lambda V^\top \\ V\Lambda U^\top & V\Lambda V^\top \end{bmatrix}$, which is PSD because it can be factorized as $[U;V]\Lambda[U;V]^\top$. More generally this is true if we multiply $(AA^\top)^{1/2}$ by λ and $(A^\top A)^{1/2}$ by $\frac{1}{\lambda}$. We therefore have

$$\left\langle \begin{bmatrix} \lambda(AA^\top)^{1/2} & A \\ A^\top & \frac{1}{\lambda}(A^\top A)^{1/2} \end{bmatrix}, \begin{bmatrix} \lambda(BB^\top)^{1/2} & -B \\ -B^\top & \frac{1}{\lambda}(B^\top B)^{1/2} \end{bmatrix} \right\rangle \ge 0, \quad (185)$$

since both terms in the inner product are PSD. This gives $\lambda^2\langle (AA^\top)^{1/2}, (BB^\top)^{1/2}\rangle + \frac{1}{\lambda^2}\langle (A^\top A)^{1/2}, (B^\top B)^{1/2}\rangle \ge 2\langle A, B\rangle$. Optimizing λ yields the claimed result. □

Next we show:

Theorem 3.6. *If A and B are matrices of the same dimensions with (sorted) lists of singular values $\sigma_1, \dots, \sigma_n$ and $\tau_1, \dots, \tau_n$, then*

$$\langle A, B\rangle \le \sum_{i=1}^n \sigma_i \tau_i. \quad (186)$$

This says that the dot product between two matrices is bounded by the dot product between their sorted singular values.

Proof. By Proposition 3.5, it suffices to show this in the case that A and B are both PSD and σ, τ are the eigenvalues. Actually we will only need A and B to be symmetric (which implies that, oddly, the inequality can hold even if some of the σ_i and τ_i are negative).

By taking similarity transforms we can assume without loss of generality that $A = \mathrm{diag}(\sigma_1, \dots, \sigma_n)$ with $\sigma_1 \ge \sigma_2 \ge \cdots \ge \sigma_n$. We thus wish to prove that $\sum_{i=1}^n \sigma_i B_{ii} \le \sum_{i=1}^n \sigma_i \tau_i$, where τ_i are the eigenvalues of B. We make use of the following lemma:

Lemma 3.7. *For all $1 \le k \le n$, we have $\sum_{i=1}^k B_{ii} \le \sum_{i=1}^k \tau_i$.*

Proof. Let B_k be the $k \times k$ top-left submatrix of B. Then $\sum_{i=1}^k B_{ii} = \mathrm{tr}(B_k)$ is the sum of the eigenvalues of B_k. We will show that the jth largest eigenvalue of B_k is smaller than the jth largest eigenvalue of B (this is a special case of the *Cauchy interlacing theorem*). We prove this using the min-max formulation of eigenvalues: $\lambda_i(M) = \min_{W:\dim(W)=i-1} \max_{v \in W^\perp, \|v\|_2 \le 1} v^\top M v$. Let W^* be the W that attains the min for $\lambda_j(B)$, and let P_k denote projection onto the first k coordinates. We have

$$\lambda_j(B_k) = \min_{W:\dim(W)=i-1} \max_{v \in W^\perp : \|v\|_2 \le 1} v^\top B_k v \tag{187}$$

$$\le \max_{v \in (W^*)^\perp : \|v\|_2 \le 1} (P_k v)^\top B_k (P_k v) \tag{188}$$

$$\le \max_{v \in (W^*)^\perp : \|v\|_2 \le 1} v^\top B v = \lambda_j(B), \tag{189}$$

which proves the lemma. □

Now with the lemma in hand we observe that, if we let $\sigma_{n+1} = 0$ for convenience, we have

$$\sum_{i=1}^n \sigma_i B_{ii} = \sum_{i=1}^n (\sigma_i - \sigma_{i+1})(B_{11} + \cdots + B_{ii}) \tag{190}$$

$$\le \sum_{i=1}^n (\sigma_i - \sigma_{i+1})(\tau_1 + \cdots + \tau_i) \tag{191}$$

$$= \sum_{i=1}^n \sigma_i \tau_i, \tag{192}$$

which yields the desired result. In the above algebra we have used *Abel summation*, which is the discrete version of integration by parts. □

Now that we have Theorem 3.6 in hand, we can easily analyze the operator and nuclear norms. Letting $\vec{\sigma}(A)$ denote the vector of non-decreasing singular values of A, we have

$$\langle Y, Z \rangle \le \langle \vec{\sigma}(Y), \vec{\sigma}(Z) \rangle \le \|\vec{\sigma}(Y)\|_\infty \|\vec{\sigma}(Z)\|_1. \tag{193}$$

This shows that the dual of the operator norm is at most the nuclear norm, since $\|\vec{\sigma}(Z)\|_1$ is the nuclear norm of Z. But we can achieve equality when $Y = U \Lambda V^\top$ by taking $Z = u_1 v_1^\top$ (then $\|Z\|_* = 1$ while $\langle Y, Z \rangle = \Lambda_{11} = \|Y\|$). So operator and nuclear norm are indeed dual to each other.

[Lecture 14]

3.1 Efficient Algorithm for Robust Regression

We now turn to the question of efficient algorithms, focusing on linear regression (we will address finite-sample issues later). Recall that information-theoretically, we found that two conditions are sufficient to imply resilience:

- *Hypercontractivity:* For all v, $\mathbb{E}_{x \sim p}[\langle x, v \rangle^4] \le \kappa \mathbb{E}_{x \sim p}[\langle x, v \rangle^2]^2$.
- *Bounded noise:* $\mathbb{E}_{x \sim p}[x z^2 x^\top] \preceq \sigma^2 \mathbb{E}_{x \sim p}[x x^\top]$.

As for mean estimation under bounded covariance, our strategy will be to check whether these two properties hold for the empirical distribution, and if they don't we will filter out points such that we guarantee removing more bad points than good points.

Unfortunately, the hypercontractivity condition is difficult to verify because it involves fourth moments. We will thus need to assume a stronger condition, called *certifiable hypercontractivity*:

$$\mathbb{E}_{x \sim p}[\langle x, v \rangle^4] \preceq_{\mathrm{sos}} \kappa \mathbb{E}_{x \sim p}[\langle x, v \rangle^2]^2, \tag{194}$$

where the LHS and RHS are considered as polynomials in v.

We will also need to introduce one additional piece of sum-of-squares machinery, called *pseudoexpectations*:

Definition 3.8. A *degree-$2k$ pseudoexpectation* is a linear map E from the space of degree-$2k$ polynomials to $\mathbb{R}$ satisfying the following two properties:

- $E[1] = 1$ (where 1 on the LHS is the constant polynomial).
- $E[p^2] \geq 0$ for all polynomials p of degree at most k.

We let $\mathcal{E}$ or $\mathcal{E}_{2k}$ denote the set of degree-$2k$ pseudoexpectations.

The space $\mathcal{E}$ can be optimized over efficiently, because it has a separation oracle expressible as a sum-of-squares program. Indeed, checking that $E \in \mathcal{E}$ amounts to solving the problem $\min\{E[p] \mid p \succeq_{\text{sos}} 0\}$, which is a sum-of-squares program because $E[p]$ is a linear function of p.

We are now ready to define our efficient algorithm for linear regression, Algorithm 4. It is closely analogous to the filter for mean estimation (Algorithm 2).

Algorithm 4 `FilterLinReg`

1: Input: $(x_1, y_1), \dots, (x_n, y_n) \in \mathbb{R}^d \times \mathbb{R}$.
2: Initialize weights $c_1, \dots, c_n = 1$.
3: Compute the empirical least squares regressor: $\hat{\theta}_c \stackrel{\text{def}}{=} (\sum_{i=1}^n c_i x_i x_i)^{-1}(\sum_{i=1}^n c_i x_i y_i)$.
4: Find, if possible, a pseudoexpectation $E \in \mathcal{E}_4$ such that $E[\frac{1}{n}\sum_{i=1}^n c_i \langle x_i, v\rangle^4] \geq 3\kappa E[(\frac{1}{n}\sum_{i=1}^n c_i \langle x_i, v\rangle^2)^2]$.
5: If E exists, let $\tau_i = E[\langle x_i, v\rangle^4]$ and update $c_i \leftarrow c_i \cdot (1 - \tau_i / \max_j \tau_j)$, and return to line 3.
6: Otherwise, find, if possible, a vector $v \in \mathbb{R}^d$ such that $\sum_{i=1}^n c_i \langle x_i, v\rangle^2 (y_i - \langle \hat{\theta}_c, x_i\rangle)^2 \geq 24\sigma^2 \sum_{i=1}^n c_i \langle x_i, v\rangle^2$.
7: If v exists, let $\tau_i = \langle x_i, v\rangle^2 (y_i - \langle \hat{\theta}_c, x_i\rangle)^2$ and update $c_i \leftarrow c_i \cdot (1 - \tau_i / \max_j \tau_j)$, and return to line 3.
8: Otherwise, output $\hat{\theta}_c$.

The algorithm first optimizes over $E \in \mathcal{E}_4$ to try to refute hypercontractivity; if it does so successfully, it filters according to $E[\langle x_i, v\rangle^4]$. Otherwise, it tries to refute the bounded noise condition, using $\hat{\theta}_c$ as a proxy for θ^* to approximate $z = y - \langle \theta^*, x\rangle$. Again, if it successfully refutes bounded noise it filters based on this. If it fails to refute either condition, we can safely output $\hat{\theta}_c$, which will be close to θ^* by resilience.

Analyzing Algorithm 4. We will show that Algorithm 4 enjoys the following loss bound:

Proposition 3.9. *Suppose that a good set S of $(1-\epsilon)n$ of the x_i satisfy:*

$$\frac{1}{n}\sum_{i\in S} \langle x_i, v\rangle^4 \preceq_{\text{sos}} \kappa(\frac{1}{n}\sum_{i\in S}\langle x_i, v\rangle^2)^2 \text{ and } \frac{1}{n}\sum_{i\in S} z_i^2 x_i x_i^\top \preceq \sigma^2 \frac{1}{n}\sum_{i\in S} x_i x_i^\top. \tag{195}$$

Then assuming $\epsilon \leq \frac{1}{100}$ and $\kappa\epsilon \leq \frac{1}{50}$, the output of Algorithm 4 has excess loss at most $250\sigma^2\epsilon$.

Proof. We analyze Algorithm 4 similarly to Algorithm 2. Specifically, we will establish the invariant that we always remove more bad points than good points. This requires showing that $\sum_{i\in S} c_i\tau_i \leq \frac{1}{2}\sum_{i=1}^n c_i\tau_i$ for both choices of τ_i in the algorithm. Concretely, we need to show:

$$\sum_{i\in S} c_i E[\langle x_i, v\rangle^4] \leq_? \frac{1}{2}\sum_{i=1}^n c_i E[\langle x_i, v\rangle^4] \text{ and } \sum_{i\in S} c_i \langle x_i, v\rangle^2 (y_i - \langle \hat{\theta}_c, x_i\rangle)^2 \leq_? \frac{1}{2}\sum_{i=1}^n c_i \langle x_i, v\rangle^2 (y_i - \langle \hat{\theta}_c, x_i\rangle)^2. \tag{196}$$

For both of these we will want the following intermediate lemma, which states that deletions of hypercontractive distributions are hypercontractive:

Lemma 3.10. *Suppose that the set S of good points is hypercontractive in the sense that $\frac{1}{n}\sum_{i\in S}\langle x_i, v\rangle^4 \preceq_{\text{sos}} \kappa(\frac{1}{n}\sum_{i\in S}\langle x_i, v\rangle^2)^2$. Then, for any c_i such that $\frac{1}{n}\sum_{i\in S}(1 - c_i) \leq \epsilon$, we have*

$$\frac{1}{n}\sum_{i\in S} c_i \langle x_i, v\rangle^4 \preceq_{\text{sos}} \frac{\kappa}{1-\kappa\epsilon}(\frac{1}{n}\sum_{i\in S} c_i \langle x_i, v\rangle^2)^2. \tag{197}$$

Proof. We expand directly; let

$$A = \frac{1}{n}\sum_{i\in S}\langle x_i, v\rangle^4, \quad B = \frac{1}{n}\sum_{i\in S}\langle x_i, v\rangle^2, \tag{198}$$

$$C = \frac{1}{n}\sum_{i\in S}(1-c_i)\langle x_i, v\rangle^4, \quad D = \frac{1}{n}\sum_{i\in S}(1-c_i)\langle x_i, v\rangle^2. \tag{199}$$

Then our goal is to show that $\frac{\kappa}{1-\kappa\epsilon}(B-D)^2 - (A-C) \succeq_{\text{sos}} 0$. We are also given that (i) $\kappa B^2 \succeq_{\text{sos}} A$ and we observe that (ii) $C \succeq_{\text{sos}} D^2/(\frac{1}{n}\sum_{i=1}^n (1-c_i) \succeq_{\text{sos}} D^2/\epsilon$ by Cauchy-Schwarz. We thus have

$$\frac{\kappa}{1-\kappa\epsilon}(B-D)^2 - (A-C) = \frac{\kappa}{1-\kappa\epsilon}B^2 - \frac{2\kappa}{1-\kappa\epsilon}BD + \frac{\kappa}{1-\kappa\epsilon}D^2 - A + C \tag{200}$$

$$\overset{(i)}{\succeq}_{\text{sos}} (\frac{\kappa}{1-\kappa\epsilon} - \kappa)B^2 - \frac{2\kappa}{1-\kappa\epsilon}BD + (\frac{\kappa}{1-\kappa\epsilon}D^2 + C) \tag{201}$$

$$\overset{(ii)}{\succeq}_{\text{sos}} (\frac{\kappa}{1-\kappa\epsilon} - \kappa)B^2 - \frac{2\kappa}{1-\kappa\epsilon}BD + (\frac{\kappa}{1-\kappa\epsilon} + \frac{1}{\epsilon})D^2 \tag{202}$$

$$= \frac{\kappa^2\epsilon}{1-\kappa\epsilon}B^2 - \frac{2\kappa}{1-\kappa\epsilon}BD + \frac{1/\epsilon}{1-\kappa\epsilon}D^2 \tag{203}$$

$$= \frac{\epsilon}{1-\kappa\epsilon}(\kappa B - D/\epsilon)^2 \succeq_{\text{sos}} 0, \tag{204}$$

as was to be shown. □

With Lemma 3.10 in hand, we proceed to analyze the filtering steps by establishing the inequalities in (196). For the first, observe that

$$\frac{1}{n}\sum_{i\in S} c_i E[\langle x_i, v\rangle^4] \overset{(i)}{\leq} \frac{\kappa}{1-\kappa\epsilon}E[(\frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^2)^2] \tag{205}$$

$$\overset{(ii)}{\leq} \frac{\kappa}{1-\kappa\epsilon}E[(\frac{1}{n}\sum_{i=1}^{n} c_i\langle x_i, v\rangle^2)^2] \tag{206}$$

$$\overset{(iii)}{\leq} \frac{1}{3(1-\kappa\epsilon)}\frac{1}{n}\sum_{i=1}^{n} c_i E[\langle x_i, v\rangle^4]. \tag{207}$$

Here (i) is by Lemma 3.10 (and the fact that $E[p] \leq E[q]$ if $p \preceq_{\text{sos}} q$), (ii) is by the fact that adding the $c_i\langle x_i, v\rangle^2$ terms for $i \notin S$ is adding a sum of squares, and (iii) is by the assumption that E refutes hypercontractivity. Thus as long as $\kappa\epsilon \leq \frac{1}{3}$ we have the desired property for the first filtering step.

For the second, observe that

$$\frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^2(y_i - \langle\hat{\theta}_c, x_i\rangle)^2 \leq \frac{2}{n}\sum_{i\in S} c_i(\underbrace{\langle x_i, v\rangle^2(y_i - \langle\theta^*, x_i\rangle)^2}_{(a)} + \underbrace{\langle x_i, v\rangle^2\langle\hat{\theta}_c - \theta^*, x_i\rangle^2}_{(b)}). \tag{208}$$

We will bound (a) and (b) in turn. To bound (a) note that

$$\frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^2(y_i - \langle\theta^*, x_i\rangle)^2 = \frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^2 z_i^2 \tag{209}$$

$$\leq \frac{1}{n}\sum_{i\in S}\langle x_i, v\rangle^2 z_i^2 \tag{210}$$

$$\leq \frac{\sigma^2}{n}\sum_{i\in S}\langle x_i, v\rangle^2 \tag{211}$$

$$\leq \frac{\sigma^2}{1-\kappa\epsilon}\frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^2, \tag{212}$$

where the last line invokes Lemma 3.10 and the middle inequality is by the bounded noise assumption for S.

To bound (b), let $R = \frac{1}{(1-\epsilon)n}\sum_{i\in S}\langle\hat{\theta}_c - \theta^*, x_i\rangle^2$, which is the excess loss of $\hat{\theta}_c$ and what we eventually hope to bound when the algorithm terminates. We use Cauchy-Schwarz and hypercontractivity:

$$\frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^2\langle\hat{\theta}_c - \theta^*, x_i\rangle^2 \le \Big(\frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^4\Big)^{1/2}\Big(\frac{1}{n}\sum_{i\in S} c_i\langle\hat{\theta}_c - \theta^*, x_i\rangle^4\Big)^{1/2} \tag{213}$$

$$\le \frac{\kappa}{1-\kappa\epsilon}\Big(\frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^2\Big)\Big(\frac{1}{n}\sum_{i\in S} c_i\langle\hat{\theta}_c - \theta^*, x_i\rangle^2\Big) \tag{214}$$

$$\le \frac{\kappa R}{1-\kappa\epsilon}\Big(\frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^2\Big). \tag{215}$$

Combining these, we obtain

$$\frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^2(y_i - \langle\hat{\theta}_c, x_i\rangle)^2 \le \frac{2\sigma^2 + 2\kappa R}{1-\kappa\epsilon}\Big(\frac{1}{n}\sum_{i\in S} c_i\langle x_i, v\rangle^2\Big). \tag{216}$$

But we are assuming that overall

$$\frac{1}{n}\sum_{i=1}^{n} c_i\langle x_i, v\rangle^2(y_i - \langle\hat{\theta}_c, x_i\rangle)^2 = S\cdot\Big(\frac{1}{n}\sum_{i=1}^{n} c_i\langle x_i, v\rangle^2\Big), \tag{217}$$

with $S \ge 10\sigma^2$. Thus we are safe as long as $\frac{2\sigma^2+2\kappa R}{1-\kappa\epsilon} \le S/2$, and the main remaining issue is to bound R in terms of S. To do so, note that the distribution weighted by c_i satisfies the hypercontractive and bounded noise conditions with parameters $\frac{3\kappa}{1-\epsilon}$ and $\frac{S}{1-\epsilon}$. It follows from Proposition 3.4 (the resilience bound for linear regression) that $R \le 10S\epsilon/(1-\epsilon)$ as long as $\epsilon(\kappa/(1-\epsilon)-1) \le \frac{1}{6}$ and $\epsilon \le \frac{1}{8}$. We thus need to verify that

$$\frac{2\sigma^2 + 20\kappa S\epsilon/(1-\epsilon)}{1-\kappa\epsilon} \le_? S/2, \tag{218}$$

which if $\epsilon\kappa \le 0.02$ and $\epsilon \le 0.01$ reduces to $2\sigma^2 + 0.4S/0.99 \le 0.49S$, which holds if $S \ge 24\sigma^2$, which is the cutoff in the algorithm.

Since the algorithm terminates with $S \le 24\sigma^2$, we incidentally also have that $R \le 250\sigma^2\epsilon$, as claimed. □

[Lectures 15-16]

4 Resilience Beyond TV Distance

We now turn our attention to distances other than the distance $D = \mathsf{TV}$ that we have considered so far. The family of distances we will consider are called *Wasserstein distances*. Given a cost function $c(x, y)$ (which is usually assumed to be a metric), we define the distance $W_c(p, q)$ between two distributions p and q as

$$W_c(p, q) = \inf_{\pi} \mathbb{E}_{x,y\sim\pi}[c(x, y)] \tag{219}$$

$$\text{subject to } \int \pi(x, y)dy = p(x), \int \pi(x, y)dx = q(y). \tag{220}$$

This definition is a bit abstruse so let us unpack it. The decision variable π is called a *coupling* between p and q, and can be thought of as a way of matching points in p with points in q ($\pi(x, y)$ is the amount of mass in $p(x)$ that is matched to $q(y)$). The Wasserstein distance is then the minimum cost coupling (i.e., minimum cost matching) between p and q. Some special cases include:

- $c(x, y) = \mathbb{I}[x \ne y]$. Then W_c is the total variation distance, with the optimal coupling being $\pi(x, x) = \min(p(x), q(x))$ (the off-diagonal $\pi(x, y)$ can be arbitrary as long as the total mass adds up correctly).
- $c(x, y) = \|x - y\|_2$. Then W_c is the *earth-mover distance*—the average amount that we need to move points around to "move" p to q.

- $c(x, y) = \|x - y\|_0$. Then W_c is the average number of coordinates we need to change to move p to q.
- $c(x, y) = \|x - y\|_2^\alpha$, for $\alpha \in [0, 1]$. This is still a metric and interpolates between TV and earthmover distance.

There are a couple key properties of Wasserstein distance we will want to use. The first is that W_c is a metric if c is:

Proposition 4.1. *Suppose that c is a metric. Then W_c is also a metric.*

Proof. TBD □

The second, called *Kantorovich-Rubinstein duality*, provides an alternate definition of W_c distance in terms of functions that are Lipschitz under c, meaning that $|f(x) - f(y)| \le c(x, y)$.

Theorem 4.2 (Kantorovich-Rubinstein)**.** *Call a function f* Lipschitz in c *if $|f(x) - f(y)| \le c(x, y)$ for all x, y, and let $\mathcal{L}(c)$ denote the space of such functions. If c is a metric, then we have*

$$W_c(p, q) = \sup_{f \in \mathcal{L}(c)} \mathbb{E}_{x\sim p}[f(x)] - \mathbb{E}_{x\sim q}[f(x)]. \tag{221}$$

As a special case, take $c(x, y) = \mathbb{I}[x \neq y]$ (corresponding to TV distance). Then $f \in \mathcal{L}(c)$ if and only if $|f(x) - f(y)| \le 1$ for all $x \neq y$. By translating f, we can equivalently take the supremum over all f mapping to $[0, 1]$. This says that

$$\mathsf{TV}(p, q) = \sup_{f:\mathcal{X}\to[0,1]} \mathbb{E}_p[f(x)] - \mathbb{E}_q[f(x)], \tag{222}$$

which recovers the definition of TV in terms of the maximum difference in probability of any event E.

As another special case, take $c(x, y) = \|x - y\|_2$. Then the supremum is over all 1-Lipschitz functions (in the usual sense).

In the next section, we will see how to generalize the definition of resilience to any Wasserstein distance.

4.1 Resilience for Wasserstein distances

We show how to extend the idea of resilience to Wasserstein distances W_c. Recall that for TV distance, we showed that resilient sets have bounded modulus $\mathfrak{m}$; this crucially relied on the midpoint property that any p_1, p_2 have a midpoint r obtained via *deletions* of p_1 or p_2. In other words, we used the fact that any TV perturbation can be decomposed into a "friendly" operation (deletion) and its opposite (addition). We think of deletion as friendlier than addition, as the latter can move the mean arbitrarily far by adding probability mass at infinity.

To extend this to other Wasserstein distances, we need to identify a similar way of decomposing a Wasserstein perturbation into a friendly perturbation and its inverse. Unfortunately, deletion is closely tied to the TV distance in particular. To get around this, we use the following re-interpretation: *Deletion is equivalent to movement towards the mean under* TV. More precisely:

> $\hat{\mu}$ is a possible mean of an ϵ-deletion of p if and only if some r with mean $\hat{\mu}$ can be obtained from p by moving points *towards* $\hat{\mu}$ with TV distance at most ϵ.

This is more easily seen in the following diagram:

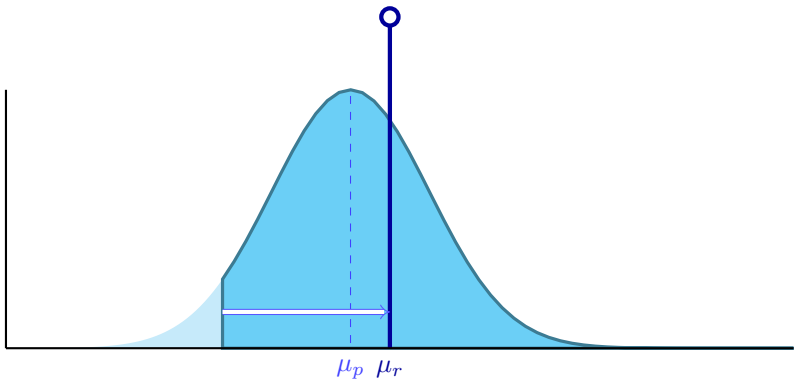

Here we can equivalently either delete the left tail of p or shift all of its mass to μ_r; both yield a modified distribution with the same mean μ_r. Thus we can more generally say that a perturbation is friendly if it only moves probability mass towards the mean. This motivates the following definition:

Definition 4.3 (Friendly perturbation)**.** For a distribution p over $\mathcal{X}$, fix a function $f : \mathcal{X} \to \mathbb{R}$. A distribution r is an ϵ-friendly perturbation of p for f under W_c if there is a coupling π between $X \sim p$ and $Y \sim r$ such that:

- The cost ($\mathbb{E}_\pi[c(X, Y)]$) is at most ϵ.
- All points move towards the mean of r: $f(Y)$ is between $f(X)$ and $\mathbb{E}_r[f(Y)]$ almost surely.

Note that friendliness is defined only in terms of one-dimensional functions $f : \mathcal{X} \to \mathbb{R}$; we will see how to handle higher-dimensional objects later. Intuitively, a friendly perturbation is a distribution r for which there exists a coupling that 'squeezes' p to μ_r.

The key property of deletion in the TV case was the existence of a *midpoint*: for any two distributions that are within ϵ in TV, one can find another distribution that is an ϵ-deletion of both distributions. We would like to show the analogous result for W_c–i.e. that if $W_c(p, q) \leq \epsilon$ then there exists an r that is an ϵ-friendly perturbation of *both* p and q for the function f.

The intuitive reason this is true is that any coupling between two one-dimensional distributions can be separated into two stages: in one stage all the mass only moves towards some point, in the other stage all the mass moves away from that point. This is illustrated in Figure 9.

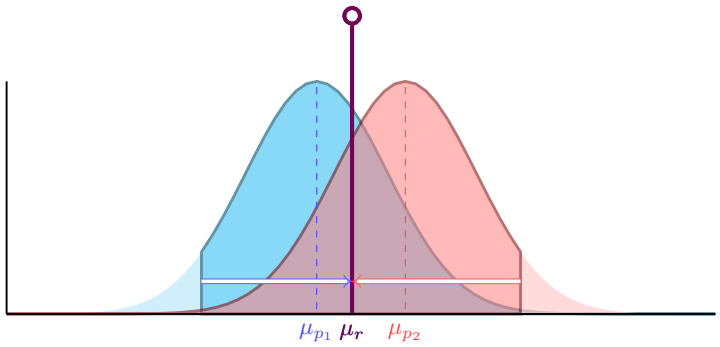

Figure 9: Illustration of midpoint lemma. For any distributions p_1, p_2 that are close under W_c, the coupling between p_1 and p_2 can be split into couplings $\pi_{p_1,r}$, $\pi_{p_2,r}$ such that p_1, p_2 only move towards μ_r under the couplings. We do this by "stopping" the movement from p_1 to p_2 at μ_r.

To formalize this intuitive argument, we need a mild topological property:

Assumption 4.4 (Intermediate value property)**.** *For any x and y and any u with $f(x) < u < f(y)$, there is some z satisfying $f(z) = u$ and $\max(c(x, z), c(z, y)) \leq c(x, y)$.*

This holds for any f if $c = \mathbb{I}[x \neq y]$ (TV distance), and for any continuous f if c is a path metric (a metric with "nice" paths between points, which includes the ℓ_2-distance). Under this assumption we can prove the desired midpoint lemma:

Lemma 4.5 (Midpoint lemma for W_c)**.** *Suppose Assumption 4.4 holds. Then for any p_1 and p_2 such that $W_c(p_1, p_2) \leq \epsilon$ and any f, there exists a distribution r that is an ϵ-friendly perturbation of both p_1 and p_2 with respect to f.*

Proof. Given any two points x and y, without loss of generality we assume $f(x) \leq f(y)$. Define

$$s_{xy}(u) = \begin{cases} \min(f(x), f(y)), & u \leq \min(f(x), f(y)) \\ u, & u \in [f(x), f(y)] \\ \max(f(x), f(y)), & u \geq \max(f(x), f(y)). \end{cases} \tag{223}$$

If we imagine u increasing from $-\infty$ to $+\infty$, we can think of s_{xy} as a "slider" that tries to be as close to u as possible while remaining between $f(x)$ and $f(y)$.

By Assumption 4.4, there must exist some point z such that $\max(c(x,z),c(z,y)) \le c(x,y)$ and $f(z) = s_{xy}(u)$. Call this point $z_{xy}(u)$.

Given a coupling $\pi(x,y)$ from p_1 to p_2, if we map y to $z_{xy}(u)$, we obtain a coupling $\pi_1(x,z)$ to some distribution $r(u)$, which by construction satisfies the squeezing property, except that it squeezes towards u rather than towards the mean $\mu(u) = \mathbb{E}_{X\sim r(u)}[f(X)]$. However, note that $u - \mu(u)$ is a continuous, monotonically non-decreasing function (since $u - s_{xy}(u)$ is non-decreasing) that ranges from $-\infty$ to $+\infty$. It follows that there is a u^* with $\mu(u^*) = u^*$. Then the couplings to $r(u^*)$ squeeze towards its mean $\mu(u^*)$.

Moreover, $\mathbb{E}_{(X,Z)\sim\pi_1}[c(X,Z)] \le \mathbb{E}_{(X,Y)\sim\pi}[c(X,Y)] = W_c(p_1,p_2)$. The coupling π_1 therefore also has small enough cost, and so is a friendly perturbation. Similarly, the coupling π_2 mapping y to $z_{xy}(u^*)$ satisfies the squeezing property and has small enough cost by the same argument. □

Defining resilience: warm-up. With Lemma 4.5 in hand, we generalize resilience to Wasserstein distances by saying that a distribution is resilient if $\mathbb{E}_r[f(X)]$ is close to $\mathbb{E}_p[f(X)]$ for every η-friendly perturbation r and every function f lying within some appropriate family $\mathcal{F}$. For now, we will focus on second moment estimation under $W_{\|\cdot\|_2}$ (we consider second moment estimation because mean estimation is trivial under $W_{\|\cdot\|_2}$). This corresponds to the loss function

$$L(p,S) = \|\mathbb{E}_{x\sim p}[xx^\top] - S\|. \tag{224}$$

For notational convenience we also typically denote $W_{\|\cdot\|_2}$ as W_1.

For the loss $L(p,S)$, we will take our family $\mathcal{F}$ to be all functions of the form $f_v(x) = \langle x,v\rangle^2$ with $\|v\|_2 = 1$. Thus we define the (ρ,ϵ)-resilient distributions under W_1 as

$$\mathcal{G}_{\text{sec}}^{W_1}(\rho,\epsilon) = \{p \mid |\mathbb{E}_r[\langle x,v\rangle^2] - \mathbb{E}_p[\langle x,v\rangle^2]| \le \rho \text{ whenever } r \text{ is } \epsilon\text{-friendly under } \langle x,v\rangle^2 \text{ and } \|v\|_2 = 1\}. \tag{225}$$

Note the twist in the definition of $\mathcal{G}_{\text{sec}}^{W_1}$–the allowed r depends on the current choice of v, since friendliness is specific to the function $f_v = \langle x,v\rangle^2$, which is different from deletions in the TV case.

We will first show that $\mathcal{G}_{\text{sec}}^{W_1}$ has small modulus, then derive sufficient moment conditions for p to be (ρ,ϵ)-resilient.

Proposition 4.6. *The set of (ρ,ϵ)-resilient distributions for W_1 has modulus $\mathfrak{m}(\mathcal{G}_{\text{sec}}^{W_1}(\rho,2\epsilon),\epsilon) \le 2\rho$.*

Proof. For a distribution q, let $S_q = \mathbb{E}_q[xx^\top]$. Suppose that $p_1,p_2 \in \mathcal{G}_{\text{sec}}^{W_1}(\rho,\epsilon)$ and $W_1(p_1,p_2) \le 2\epsilon$. For any v, by Lemma 4.5, there exists an r that is a (2ϵ)-friendly perturbation of both p_1 and p_2 with respect to $\langle x,v\rangle^2$. We conclude that $|\mathbb{E}_{p_i}[\langle x,v\rangle^2] - \mathbb{E}_r[\langle x,v\rangle^2]| \le \rho$ for $i = 1,2$, and hence $|\mathbb{E}_{p_1}[\langle x,v\rangle^2] - \mathbb{E}_{p_2}[\langle x,v\rangle^2]| \le 2\rho$, which can be written as $|v^\top(S_{p_1} - S_{p_2})v| \le 2\rho$. Taking the sup over $\|v\|_2 = 1$ yields $\|S_{p_1} - S_{p_2}\| \le 2\rho$. Since $L(p_1,\theta^*(p_2)) = \|S_{p_1} - S_{p_2}\|$, this gives the desired modulus bound. □

Sufficient conditions for W_1-resilience. Recall that for mean estimation under TV perturbation, any distribution with bounded ψ-norm was $(\mathcal{O}(\epsilon\psi^{-1}(1/\epsilon)),\epsilon)$-resilient. In particular, bounded covariance distributions were $(\mathcal{O}(\sqrt{\epsilon}),\epsilon)$-resilient. We have an analogous result for W_1-resilience, but with a modified ψ function:

Proposition 4.7. *Let ψ be an Orlicz function, and define $\tilde{\psi}(x) = x\psi(2x)$. Suppose that X (not $X-\mu$) has bounded $\tilde{\psi}$-norm: $\mathbb{E}_p[\tilde{\psi}(|v^\top X|/\sigma)] \le 1$ for all unit vectors v. Also assume that the second moment of p is at most σ^2. Then p is (ρ,ϵ) resilient for $\rho = \max(\sigma\epsilon\psi^{-1}(\frac{2\sigma}{\epsilon}), 4\epsilon^2 + 2\epsilon\sigma)$.*

Let us interpret Proposition 4.7 before giving the proof. Take for instance $\psi(x) = x^2$. Then Proposition 4.7 asks for the 3rd moment to be bounded by $\sigma^3/4$. In that case we have $\rho = \sigma\epsilon\psi^{-1}(2\sigma/\epsilon) = \sqrt{2}\sigma^{3/2}\epsilon^{1/2}$. If the units seem weird, remember that ϵ has units of distance (before it was unitless) and hence $\sigma^{3/2}\epsilon^{1/2}$ has quadratic units, which matches the second moment estimation task.

More generally, taking $\psi(x) = x^k$, we ask for a $(k+1)$st moment bound and get error $\mathcal{O}(\sigma^{1+1/k}\epsilon^{1-1/k})$.

We now turn to proving Proposition 4.7. A helpful auxiliary lemma (here and later) proves a way to use Orlicz norm bounds:

Lemma 4.8. *Let p and q be two distributions over $\mathcal{X}$, $g : \mathcal{X} \to \mathbb{R}$ be any function, c be a non-negative cost function, and ψ be an Orlicz function. Then for any coupling $\pi_{p,q}$ between p and q and any $\sigma > 0$ we have*

$$|\mathbb{E}_{X\sim p}[g(X)] - \mathbb{E}_{Y\sim q}[g(Y)]| \le \sigma \mathbb{E}_{\pi_{p,q}}[c(X,Y)]\psi^{-1}\left(\frac{\mathbb{E}_{\pi_{p,q}}\left[c(X,Y)\psi\left(\frac{|g(X)-g(Y)|}{\sigma c(X,Y)}\right)\right]}{\mathbb{E}_{\pi_{p,q}}[c(X,Y)]}\right). \tag{226}$$

Proof. Note that $|\mathbb{E}_p[g(X)] - \mathbb{E}_q[g(Y)]| = |\mathbb{E}_\pi[g(X) - g(Y)]|$. We weight the coupling π by the cost c to obtain a new probability measure $\pi'(x,y) = c(x,y)\pi(x,y)/\mathbb{E}[c(x,y)]$. We apply Jensen's inequality under π' as follows:

$$\psi\Big(\Big|\frac{\mathbb{E}_\pi[g(X)-g(Y)]}{\sigma\mathbb{E}_\pi[c(X,Y)]}\Big|\Big) = \psi\Big(\Big|\mathbb{E}_\pi\Big[\frac{c(X,Y)}{\mathbb{E}[c(X,Y)]}\cdot\frac{g(X)-g(Y)}{\sigma c(X,Y)}\Big]\Big|\Big) \tag{227}$$

$$= \psi\Big(\Big|\mathbb{E}_{\pi'}\Big[\frac{g(X)-g(Y)}{\sigma c(X,Y)}\Big]\Big|\Big) \tag{228}$$

$$\le \mathbb{E}_{\pi'}\Big[\psi\Big(\frac{|g(X)-g(Y)|}{\sigma c(X,Y)}\Big)\Big] \tag{229}$$

$$= \mathbb{E}_\pi\Big[c(X,Y)\psi\Big(\frac{|g(X)-g(Y)|}{\sigma c(X,Y)}\Big)\Big]/\mathbb{E}_\pi[c(X,Y)]. \tag{230}$$

Inverting ψ yields the desired result. □

Proof of Proposition 4.7. We apply Lemma 4.8 with $q = r$ an ϵ-friendly perturbation of p under $\langle x,v\rangle^2$, and $g = \langle x,v\rangle^2$; we will also use cost $c'(x,y) = |v^\top(x-y)|$, which satisfies $c'(x,y) \le c(x,y)$. Taking π to be the ϵ-friendly coupling (under c, not c') between p and r yields

$$|\mathbb{E}_p[\langle x,v\rangle^2] - \mathbb{E}_r[\langle x,v\rangle^2]| \le \sigma\epsilon\psi^{-1}\left(\frac{\mathbb{E}_\pi\left[|\langle x-y,v\rangle|\psi\left(\frac{|\langle x,v\rangle^2-\langle y,v\rangle^2|}{\sigma|\langle x-y,v\rangle|}\right)\right]}{\epsilon}\right) \tag{231}$$

$$= \sigma\epsilon\psi^{-1}\left(\frac{\mathbb{E}_\pi\left[|\langle x-y,v\rangle|\psi\left(|\langle x,v\rangle + \langle y,v\rangle|/\sigma\right)\right]}{\epsilon}\right). \tag{232}$$

Now we will split into two cases. First, we observe that the worst-case friendly perturbation will either move all of the $\langle x,v\rangle^2$ upwards, or all of the $\langle x,v\rangle^2$ downwards, since otherwise we could take just the upwards part or just the downwards part and perturb the mean further. In other words, we either have (i) $\langle x,v\rangle^2 \ge \langle y,v\rangle^2$ for all $(x,y) \in \mathrm{supp}(\pi)$ with $x \ne y$, or (ii) $\langle x,v\rangle^2 \le \langle y,v\rangle^2$ for all $(x,y) \in \mathrm{supp}(\pi)$ with $x \ne y$. We analyze each case in turn.

Case (i): y moves downwards. In this case we can use the bounds $|\langle x-y,v\rangle| \le 2|\langle x,v\rangle|$ and $|\langle x+y,v\rangle| \le 2|\langle x,v\rangle|$ together with (232) to conclude that

$$|\mathbb{E}_p[\langle x,v\rangle^2] - \mathbb{E}_r[\langle x,v\rangle^2]| \le \sigma\epsilon\psi^{-1}\Big(\mathbb{E}_\pi\Big[2|\langle x,v\rangle|\psi\Big(\frac{2|\langle x,v\rangle|}{\sigma}\Big)\Big]/\epsilon\Big) \tag{233}$$

$$= \sigma\epsilon\psi^{-1}\Big(\mathbb{E}_p\Big[2\sigma\tilde{\psi}\Big(\frac{|\langle x,r\rangle|}{\sigma}\Big)\Big]/\epsilon\Big) \tag{234}$$

$$\le \sigma\epsilon\psi^{-1}(2\sigma/\epsilon), \tag{235}$$

where the final inequality is by bounded Orlicz norm of p.

Case (ii): y moved upwards. In this case by friendliness we have that $|\langle y,v\rangle|^2 \le v^\top S_r v$ whenever $(x,y) \in \mathrm{supp}(\pi)$ and $y \ne x$. Thus

$$|\langle x-y,v\rangle|\psi(|\langle x,v\rangle + \langle y,v\rangle|/\sigma) \le |\langle x-y,v\rangle|\psi(2|\langle y,v\rangle|/\sigma) \le |\langle x-y,v\rangle|\psi(2\sqrt{v^\top S_r v}/\sigma). \tag{236}$$

for all $(x,y) \in \mathrm{supp}(\pi)$. Plugging back into (232) yields

$$|\mathbb{E}_p[\langle x,v\rangle^2] - \mathbb{E}_r[\langle x,v\rangle^2]| \le \sigma\epsilon\psi^{-1}(\mathbb{E}_\pi[|\langle x-y,v\rangle|\psi(2\sqrt{v^\top S_r v}/\sigma)]/\epsilon) \tag{237}$$

$$\le \sigma\epsilon\psi^{-1}(\epsilon\cdot\psi(2\sqrt{v^\top S_r v}/\sigma)/\epsilon) \tag{238}$$

$$= \sigma\epsilon\cdot 2v^\top S_r v/\sigma = 2\epsilon\sqrt{v^\top S_r v}. \tag{239}$$

Here the final inequality is because $\mathbb{E}_\pi[|\langle x-y,v\rangle|] \le \mathbb{E}_\pi[c(x,y)] \le \epsilon$ under the coupling. Comparing the left-hand-side to the final right-hand-side yields $|v^\top S_p v - v^\top S_r v| \le 2\epsilon\sqrt{v^\top S_r v}$. Thus defining $\Delta = |v^\top S_p v - v^\top S_r v|$ and using the fact that $v^\top S_p v \le \sigma^2$, we obtain $\Delta \le 2\epsilon\sqrt{\Delta+\sigma^2}$, which implies (after solving the quadratic) that $\Delta \le 4\epsilon^2 + 2\epsilon\sigma$.

Thus overall we have $|\mathbb{E}_p[\langle x,v\rangle^2] - \mathbb{E}_r[\langle x,v\rangle^2]| \le \max(\sigma\epsilon\psi^{-1}(2\sigma/\epsilon), 4\epsilon^2 + 2\epsilon\sigma)$, as was to be shown. □

4.2 Other Results

Our understanding of robust estimation under W_c distances is still rudimentary. Below are a couple of known results, but many of these may be improved or extended in the near future (perhaps by you!).

The most straightforward extension is from second moment estimation to kth moment estimation. In that case instead of using $\tilde{\psi}(x) = x\psi(2x)$, we use $\tilde{\psi}(x) = x\psi(kx^{k-1})$. Essentially the same proof goes through.

We can also extend to more general loss functions $L(p,\theta)$, as long as L is a convex function of p for fixed θ (this holds e.g. for any $L(p,\theta) = \mathbb{E}_{x\sim p}[\ell(\theta;x)]$, since these loss functions are linear in p and hence also convex). Here the main challenge is defining an appropriate family $\mathcal{F}$ of functions for which to consider friendly perturbations. For second moment estimation our family $\mathcal{F}$ was motivated by the obsevation that $L(p,S) = \sup\{|\mathbb{E}_p[f_v(x)] - \mathbb{E}_q[f_v(x)]| \mid f_v(x) = \langle x,v\rangle^2, \|v\|_2 = 1\}$, but such linear structure need not hold in general. But we can still exploit linear structure by looking at subgradients of the loss. In particular, we can take the Fenchel-Moreau representation

$$L(p,\theta) = \sup_{f\in\mathcal{F}_\theta} \mathbb{E}_{x\sim p}[f(x)] - L^*(f,\theta), \tag{240}$$

which exists for some $\mathcal{F}_\theta$ and L^* whenever $L(p,\theta)$ is convex in p. The family $\mathcal{F}_\theta$ is roughly the family of subgradients of $L(p,\theta)$ as p varies for fixed θ. In this case we obtain conditions $G_\downarrow$ and $G_\uparrow$ as before, asking that

$$\mathbb{E}_r[f(x)] - L^*(f,\theta^*(p)) \le \rho_1 \text{ for all } f \in \mathcal{F}_{\theta^*(p)} \text{ and } \epsilon\text{-friendly } r, \tag{$\downarrow$}$$

and furthermore

$$L(p,\theta) \le \rho_2 \text{ if for every } f\in\mathcal{F}_\theta \text{ there is an } \epsilon\text{-friendly } r \text{ such that } \mathbb{E}_r[f(x)] - L^*(f,\theta) \le \rho_1. \tag{$\uparrow$}$$

Note that for the second condition ($\mathcal{G}_\downarrow$), we allow the perturbation r to depend on the current function f. If r was fixed this would closely match the old definition, but we can only do that for deletions since in general even the set of feasible r depends on f.

Using this, we can (after sufficient algebra) derive sufficient conditions for robust linear regression under W_1, for conditions similar to the hypercontractivity condition from before. This will be a challenge problem on the homework.

Finally, we can define a $\tilde{W}_1$ similar to $\tilde{\mathsf{TV}}$, but our understanding of it is far more rudimentary. In particular, known analyses do not seem to yield the correct finite-sample rates (for instance, the rate of convergence includes an $n^{-1/3}$ term that seems unlikely to actually exist).

[Lecture 17]

5 Test-Time Robustness

We now consider a different setting, focused on modeling corruptions at training time, rather than test time. For now assume that we are solving a classification task, so we wish to a predict a label $y \in \mathcal{Y}$ given covariates x (for instance, $\mathcal{Y} = \{-1,+1\}$ corresponds to binary classification, or $\mathcal{Y} = \{0,\dots,k-1\}$ to k-class classification). For a population distribution p, the setting is as follows:

- At train time, we observe samples $(x_1,\dots,x_n) \sim p$.
- A test sample is generated by first drawing $(x,y)\sim p$, and replacing x with $\bar{x}$ such that $d(x,\bar{x}) \le \epsilon$ for some distance d. We then observe $\bar{x}$ and hope to output y.

This can be placed into our previous framework by letting $\tilde{p} = p$ and saying that $D(p^*, \tilde{p}) \leq \epsilon$ if there is a coupling π from $\tilde{p}$ to p^* such that $d(x, \bar{x}) \leq \epsilon$ and $y = y'$ almost surely for $(x, y, \bar{x}, y') \sim \pi$. In the language of Wasserstein distance, this says that $W_{d,\infty}(\tilde{p}, p^*) \leq \epsilon$ (here we define $W_{c,k}(p, q) = \inf_{\pi \in \Pi(p,q)} \mathbb{E}_{(z,z') \sim \pi}[c(z, z')^k]^{1/k}$).

However, there are some important differences from the setting from before:

- We generally think of $p = \tilde{p}$ as the "nice" distribution and p^* as the "ugly" distribution.
- Since p^* is the "ugly" distribution, we make no distributional assumptions $\mathcal{G}$ about p^*.
- On the other hand, we also generally consider much "smaller" perturbations than in the previous case. For instance, $d(x, \bar{x}) \leq \epsilon$ should imply that x and $\bar{x}$ are close enough that they have the same label y.

A final point is that the worst-case test loss is directly observable, at least given infinitely many samples from p. Indeed, if the (non-robust) loss is

$$L_0(p, \theta) = \mathbb{E}_{(x,y) \sim p}[\ell(\theta; x, y)], \tag{241}$$

then the robust loss is

$$L(p, \theta) = \mathbb{E}_{(x,y) \sim p}\left[\sup_{\bar{x}: d(x,\bar{x}) \leq \epsilon} \ell(\theta; \bar{x}, y)\right]. \tag{242}$$

(The worst-case test loss was also observable for train-time corruptions, but it involved a much more complex supremum and depended on $\mathcal{G}$.)

Consequently, given samples $(x_1, y_1), \ldots, (x_n, y_n)$, a natural estimator is

$$\hat{\theta} = \arg\min_{\theta} \rho(\theta) + \frac{1}{n} \sum_{i=1}^{n} \sup_{d(x_i, \bar{x}_i) \leq \epsilon} \ell(\theta; \bar{x}_i, y_i), \tag{243}$$

where $\rho(\theta)$ is a regularizer and the other term is an empirical estimate of the worst-case loss. Indeed, this is essentially the estimator we will consider throughout. However, there are several issues to address:

- How well does $\hat{\theta}$ generalize from the training points to the population distribution?
- The supremum over $d(x, \bar{x}) \leq \epsilon$ is generally computationally intractable. What happens if we only approximate this supremum, and what are good approximation schemes?
- How robust are results to the choice of d itself?

Answering these questions will involve a mix of empirical data and theoretical analysis.

5.1 Getting Oriented: Examples of d and basic empirical results

A basic motivating example takes x to be an image containing some object y. Thus for instance $x \in [0, 1]^{224 \times 224 \times 3}$ is a 150528-dimensional vector (corresponding to a 224×224 RGB image), and y is one of k different object classes (in the most popular computer vision benchmark, ImageNet, we have $k = 1000$). We then take $d(x, \bar{x}) = \|x - \bar{x}\|_\infty$, meaning that we are allowed to perturb each pixel by at most ϵ. A striking phenomenon is that regularly-trained machine learning models are often non-robust at even very small values of ϵ. For instance, the following two images have different classifications under a neural network:

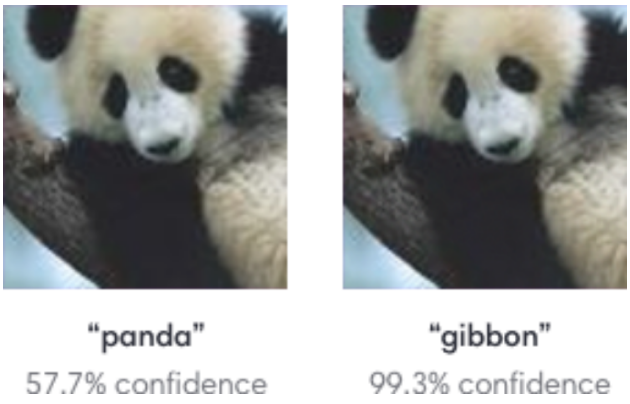

This corresponds to a perturbation in ℓ_∞ of $\epsilon = 0.007$. This is not an isolated phenomenon, and many machine learning models can be successfully adversarially perturbed at small ϵ for most of their inputs.

We could consider many perturbation measures beyond ℓ_∞:

- $d(x, \bar{x}) = \|x\bar{x}\|_\infty$: the maximum amount any pixel changes
- $d(x, \bar{x}) = \|x - \bar{x}\|_1$: the total amount that the pixels change
- $d(x, \bar{x}) = \|\mathsf{JPEG}(x) - \mathsf{JPEG}(\bar{x})\|_2$: the ℓ_2-distance in wavelet space
- The amount of "stretch" needed to move x to $\bar{x}$
- The amount of "fog", "snow", etc. needed to move x to $\bar{x}$.

These and others are illustrated below:

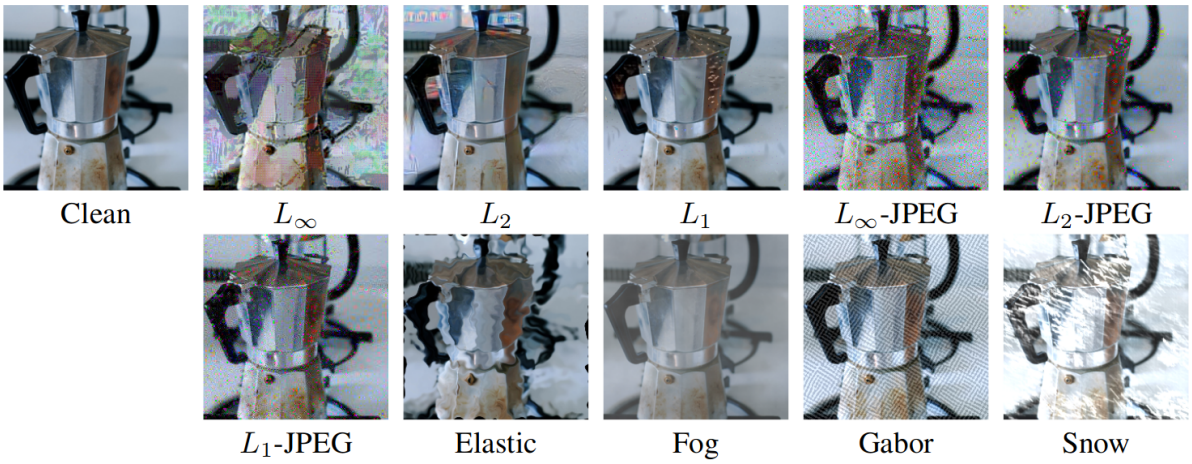

Note that for many of these, d is not a metric (it isn't even symmetric, e.g. for snow).

There are four important phenomena that we will elaborate on below:

- Adversarial perturbations have a much larger effect than random perturbations.
- Adversarial robustness is difficult to measure and most papers get it wrong enough that the results are meaningless.
- Robustness for one choice of d only partially generalizes to other choices of d.
- Minimizing the robust loss seems to increase overfitting for neural networks.

Worst-case vs. random perturbations. As a simplified case suppose that we are doing binary classification and our model is linear: $f_\theta(x) = \text{sign}(\langle \theta, x \rangle)$. We then consider two types of perturbations:

- Random Gaussian perturbations: $\bar{x} = x + z$, where $z \sim \mathcal{N}(0, \sigma^2 I)$ (standard deviation σ per coordinate).
- Worst-case ℓ_2-perturbations: $\bar{x} = x + z$, where $\|z\|_2 \leq \sigma\sqrt{d}$.

Note that the ℓ_2-norm of the perturbation is approximately the same in both cases. For a point x, define the *margin* $\gamma_x = |\langle \theta, x \rangle| / \|\theta\|_2$. How big does the margin need to be for the output f_θ to be robust to perturbations?

For the Gaussian perturbations, we have $\langle \theta, \bar{x} \rangle = \langle \theta, x \rangle + \sigma \|\theta\|_2 \mathcal{N}(0, 1)$, so the output is robust as long as $\gamma_x \gg \sigma$. However, for worst-case ℓ_2-perturbations we have $\langle \theta, \bar{x} \rangle = \langle \theta, x \rangle + \langle \theta, z \rangle$, which in the worst-case is $\langle \theta, x \rangle \pm \sigma\sqrt{d}\|\theta\|_2$. Thus we are only robust to worst-case perturbations when $\gamma_x > \sigma\sqrt{d}$. There is a $\sqrt{d}$ factor difference between random and worst-case perturbations! In fact, a Gaussian perturbation with *per-coordinate* magnitude σ is most similar to a worst-case perturbation with *overall* magnitude σ.

Adversarial robustness is difficult to measure. To measure adversarial robustness, we seek to compute

$$\ell_{\text{robust}}(\theta; x, y) = \sup_{\bar{x}: d(x,\bar{x}) \leq \epsilon} \ell(\theta; \bar{x}, y) \tag{244}$$

for a given x and y, which is computationally intractable in general. For instance, ℓ is often the loss function for a neural network, which is highly non-convex not just in θ but also in x (which is the relevant variable for our purposes). One idea would be to approximate the supremum via some local search algorithm (e.g. projected gradient descent). This gives us some proxy loss $\ell_{\text{proxy}} \leq \ell_{\text{robust}}$.

The issue arises when we take the natural step of optimizing ℓ_{proxy} to try to find a robust model θ. Since ℓ_{proxy} underestimates ℓ_{robust}, this highly incentivizes finding θ where ℓ_{proxy} is a bad approximation to ℓ_{robust}! For instance, if our proxy loss is obtained via gradient descent, it incentivizes θ for which gradient descent on $\bar{x}$ doesn't converge well. Even if we don't explicitly optimize ℓ_{proxy}, if we design models with ℓ_{proxy} in mind we as the researcher might accidentally find ways to make ℓ_{proxy} but not ℓ_{robust} small. And since ℓ_{proxy} is the only thing we measure, we have no straightforward way of even noticing this!

I estimate that around 95% of papers on test-time robustness succumb to this issue, where the proxy loss measured in the paper is small but the true robust loss is large. Some good papers that discuss this are Athalye et al. (2018) and Carlini et al. (2019), which both also provide best practices for evaluating robustness. (See also Carlini and Wagner (2017) for an earlier paper making similar points for the detection of adversarial perturbations.) Some example best practices are:

- If your proxy is based on some optimization algorithm, give the optimizer much more power at evaluation time than at training time (e.g. use 10 gradient descent steps at train time but 1000 at test time).
- Ensure that your proxy works well by making sure that when ϵ is large enough that accuracy should be zero, the proxy also gives an accuracy of zero.
- Make sure that other correlated measures (such as robustness to random perturbations) are also improving. This is more controversial because not everyone agrees on what measures should correlate with adversarial robustness.

Robustness to choice of d. While this is a nuanced story, the zeroth order summary is that robust for fairly different choices of d is positively correlated at low ϵ, but eventually becomes negatively correlated at high ϵ (although aiming for robustness under the wrong d often still gives better results than doing nothing). Here is a grid illustrating this:

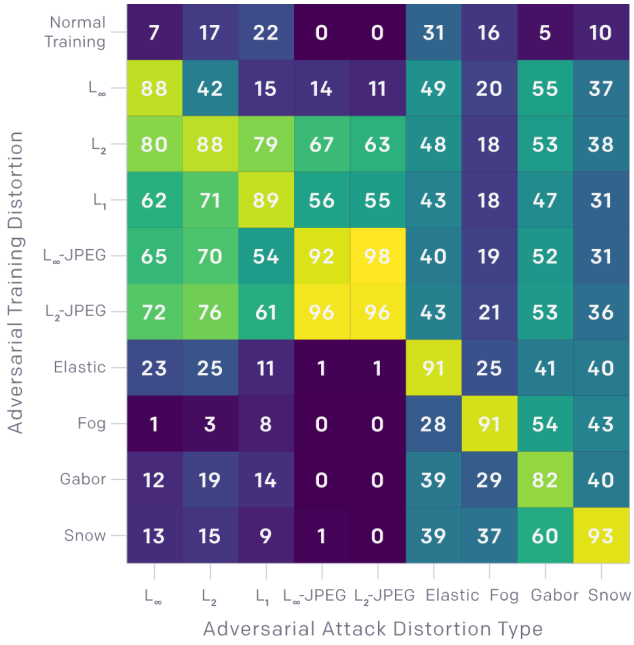

Here each row is a model that was trained to be robust to a certain type of perturbation, and each column is a measure of the robustness against a perturbation (higher is better). We will discuss in the next section how to actually perform this training. We see that even for fairly perturbations as different as elastic warping and snow, training against one perturbation helps with the other one. See Kang et al. (2019) for a more detailed investigation of this.

Robustness hurts generalization. Training a model to be robust, especially at large ϵ, can harm generalization performance. This becomes even more pronounced if we allow multiple perturbations, for instance if d allows both elastic warping and ℓ_∞ perturbations, we see overfitting even at a moderately-sized $\epsilon = 4$ (this is 4/255 so approximately 0.016):

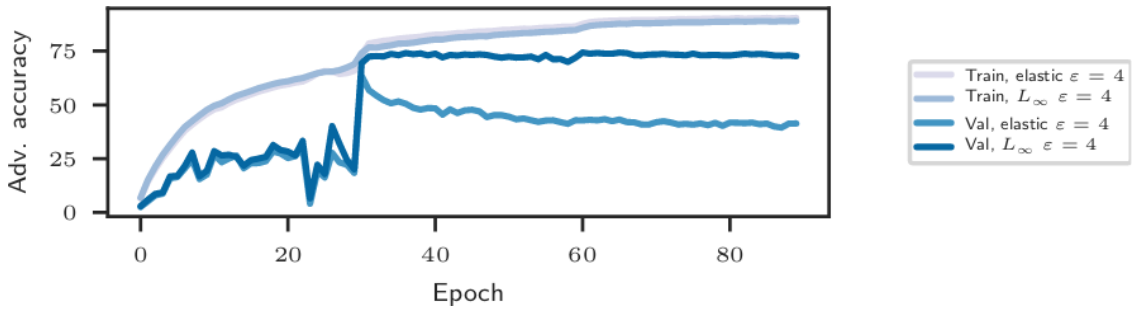

We see similar phenomena for a single perturbation, but need ϵ to be larger (e.g. around 32/255 for ℓ_∞, which is large enough to be hard for humans).

[Lecture 18]

5.2 Adversarial Training

We next discuss how to minimize the robust loss, which we recall is

$$L(p, \theta) = \mathbb{E}_{(x,y)\sim p}\left[\max_{\bar{x}: d(x,\bar{x})\leq\epsilon} \ell(\theta; \bar{x}, y)\right] \tag{245}$$

Our focus will be on computing the gradient $\nabla_\theta L(p, \theta)$ with respect to θ. If we can compute this gradient, then we can typically minimize $L(p, \theta)$ by gradient descent.

The *envelope theorem* show us how to compute the gradient of a maximum:

Theorem 5.1 (Envelope theorem)**.** *Suppose that $f(\theta, x)$ is continuously differentiable in θ and x, and that $\mathcal{B}$ is a compact set. Then*

$$\nabla_\theta \max_{x\in\mathcal{B}} f(\theta, x) = \nabla_\theta f(\theta, x)|_{x=\arg\max_{x\in\mathcal{B}} f(\theta,x)}, \tag{246}$$

assuming the arg max *is unique. In other words, the gradient of a max is the gradient at the argmax.*

This is actually a special case; the envelope theorem also covers the case where the constraints can depend on θ. Additionally, we can often relax the compactness condition on $\mathcal{B}$, although some topological condition is needed. Finally, in the convex case Danskin's theorem handles situations where the arg max is not unique, and instead shows that the subgradient is the convex hull of $\{\nabla_\theta f(\theta, x) \mid x \in \arg\max\{f(\theta, x) \mid x \in \mathcal{B}\}\}$.

The envelope theorem motivates the following procedure, called **adversarial training**. Given training points $(x_1, y_1), \ldots, (x_n, y_n)$, we do the following:

- Sample a batch of examples $\{(x_i, y_i) \mid i \in S\}$.
- For each example in the batch, let $\bar{x}_i$ approximately maximize $\ell(\theta; \bar{x}, y_i)$ such that $d(x_i, \bar{x}) \leq \epsilon$.
- Update θ in the direction $\frac{1}{|S|}\sum_{i\in S} \nabla_\theta \ell(\theta; \bar{x}_i, y_i)$.

If there is a regularizer, we would include the regularizer in the update as well.

If we exactly maximize over $\bar{x}$, then this exactly computes a (stochastic) gradient of the robust loss with respect to θ. If we only approximately maximize over $\bar{x}$, then this procedure has no guarantees, but we might hope that sufficiently good approximations lead to a good value of θ (this is only sometimes the case).

As mentioned in the previous section, a typical way of approximating the maximization would be projected gradient ascent in $\bar{x}$: take a gradient step in $\bar{x}$, then project onto the set $\{\bar{x} \mid d(x_i, \bar{x}) \leq \epsilon\}$. There isn't much to say about this beyond the empirical story, so we will focus on an alternative procedure called *certified adversarial training*, based on minimizing an upper bound on the robust loss.

5.3 Certified Adversarial Training

Define $\ell_{\mathsf{robust}}(\theta, x) = \max_{\bar{x}: d(x,\bar{x}) \leq \epsilon} \ell(\theta; \bar{x}, y)$. In certified adversarial training, we seek to construct a function ℓ_{cert} such that $\ell_{\mathsf{cert}}(\theta, x, y) \geq \ell_{\mathsf{robust}}(\theta, x, y)$ for all x, θ. Then we will minimize $\mathbb{E}[\ell_{\mathsf{cert}}(\theta; x, y)]$. Minimizing this upper bound is more desirable than minimizing a lower bound, since ℓ_{cert} is large when it is a bad approximation to ℓ_{robust}, and so the optimizer will tend to avoid those regions and find regions where ℓ_{cert} (and hence ℓ_{robust}) is small.

We will consider two strategies for constructing an appropriate ℓ_{cert}. Both are based on convex relaxations; the first is based on an LP relaxation and the second is based on an SDP relaxation. In both cases, we will need to make use of the specific structure of the function ℓ. Here we assume that ℓ corresponds to a neural network, and so can be obtained via the following system of equations:

$$\begin{aligned}
x^{(0)} &= x \\
z^{(1)} &= A^{(0)} x^{(0)} \\
x^{(1)} &= \max(z^{(1)}, 0) \\
z^{(2)} &= A^{(1)} x^{(1)} \\
&\vdots \\
x^{(L-1)} &= \max(z^{(L-1)}, 0) \\
z^{(L)} &= A^{(L-1)} x^{(L-1)} \\
\ell &= \log(\sum_j \exp(z_j^{(L)})) - z_y^{(L)}
\end{aligned} \tag{247}$$

Here the $z^{(i)}$ and $x^{(i)}$ are pre- and post-activation firings in the neural network, and the final loss ℓ is the cross-entropy loss, which is the negative log-probability assigned to y under the distribution $\pi(y) \propto \exp(z_y^{(L)})$. Thus in particular the dimension of $x^{(0)}$ is the same dimension d as that of x, while the dimension of $z^{(L)}$ is the number of classes k.

For simplicity we will consider the two-class case, where up to a monotone transformation we wish to bound $z_2^{(L)} - z_1^{(L)}$. We will also specialize to the distance $d(x, \bar{x}) = \|x - \bar{x}\|_\infty$. Thus, ℓ_{cert} can be obtained from any valid upper bound on the optimization problem

$$\underset{x^{(0:L-1)}, z^{(1:L)}}{\text{maximize}} \quad z_2^{(L)} - z_1^{(L)} \tag{248}$$

$$\text{subject to } z^{(i+1)} = A^{(i)} x^{(i)} \qquad \text{for } i = 0, \ldots, L-1 \tag{249}$$

$$x^{(i)} = \max(z^{(i)}, 0) \qquad \text{for } i = 0, \ldots, L-1 \tag{250}$$

$$|x_j^{(0)} - x_j| \leq \epsilon \qquad \text{for } j = 1, \ldots, d. \tag{251}$$

The first and last constraints are both linear inequality constraints ($|z| \leq \epsilon$ can be written as $z \leq \epsilon$ and $z \geq -\epsilon$), so the only source of non-convexity is the constraint $x^{(i)} = \max(z^{(i)}, 0)$. The LP and SDP relaxations are two different ways of handling this constraint.

Linear programming relaxation. The most obvious way to relax $x = \max(z, 0)$ to an LP is to replace it with $x \geq \max(z, 0)$, which corresponds to the two linear constraints $x \geq z$, $x \geq 0$. However, these constraints

allow x to be infinitely large, which in almost all cases will make the LP itself infinite as well (since e.g. we can make any coordinates of $x^{(L-1)}$ infinitely large, which will allow $z_2^{(L)} - z_1^{(L)}$ to become infinitely large if $A_2^{(L-1)} - A_1^{(L-1)}$ has any positive entries).

To fix this, suppose for a moment that we know that $z \in [l, u]$ for some lower and upper bounds $l \le 0 \le u$. Then we also know that $x \le \frac{z-l}{u-l}u$ (this takes on the value 0 at $z = l$ and u at $z = u$). Thus we can relax $x = \max(z, 0)$ to the three constraints:

$$x \ge z, \quad x \ge 0, \quad x \le \frac{z-l}{u-l}u. \tag{252}$$

Now x is bounded both above and below under the constraints, so we can obtain non-trivial bounds. Also observe that if our bounds $[l, u]$ instead satisfied $l \le u \le 0$ or $0 \le l \le u$ (i.e. l and u have the same sign) then the constraint is even simpler since we know that either $x = 0$ (if $l, u \le 0$) or $x = z$ (if $l, u \ge 0$). Thus in all cases we achieve some sort of bound on x, and all that remains is to explain how to compute l and u.

Computing coordinate-wise bounds. Conceptually, we can compute the bounds l and u inductively from layer to layer. To start with, we have $l^{(0)} \le x^{(0)} \le u^{(0)}$ where $l_j^{(0)} = -\epsilon$, $u_j^{(0)} = \epsilon$ for all j; this is just the ℓ_∞ constraint. Then for computing $u^{(I)}$ assuming we already know $l^{(0:I-1)}, u^{(0:I-1)}$, we solve the following LP for each coordinate j:

$$\underset{x^{(0:I-1)}, z^{(1:I)}}{\text{maximize}} \; z_j^{(I)} \tag{253}$$

$$\text{subject to } z^{(i+1)} = A^{(i)}x^{(i)} \qquad \text{for } i = 0, \dots, I-1 \tag{254}$$

$$x^{(i)} \ge z^{(i)}, \quad x^{(i)} \ge 0, \quad x^{(i)} \le \frac{z^{(i)} - l^{(i)}}{u^{(i)} - l^{(i)}}u^{(i)} \qquad \text{for } i = 0, \dots, I-1 \tag{255}$$

$$-\epsilon \le x^{(0)} \le \epsilon \tag{256}$$

We can compute $l^{(I)}$ similarly by minimizing instead of maximizing. Of course this is very inefficient since it requires solving a linear program for every single node in the network (and every single training example). We will discuss below how to avoid these costs. Note that these are both of the same form as the original program, but with a different objective instead of the original $z_2^{(L)} - z_1^{(L)}$.

Faster coordinate-wise bounds via duality. To obtain bounds, it is helpful to work in the dual, since any feasible setting of the dual variables yields an upper bound. Define dual variables $\lambda^{(i+1)}$ for the equality constraint, $\lambda_+^{(i)}$, $\mu_+^{(i)}$ for the $x \ge z$ and $x \ge 0$ constraints, and $\lambda_-^{(i)}$ for the linear upper bound constraint on x. Also define variables η_+, η_- for the $x^{(0)} \ge x - \epsilon$ and $x^{(0)} \le x + \epsilon$ constraints. Then the dual becomes

$$\underset{\lambda, \lambda_+ \ge 0, \mu_+ \ge 0, \lambda_- \ge 0, \eta_+ \ge 0, \eta_- \ge 0}{\text{minimize}} \; \sum_{i=1}^{I-1} \langle \lambda_-^{(i)}, \frac{u^{(i)}l^{(i)}}{u^{(i)} - l^{(i)}} \rangle + \epsilon \mathbb{1}^\top (\eta_+ + \eta_-) \tag{257}$$

$$\text{subject to } \lambda_-^{(i)} - \lambda_+^{(i)} - \mu_+^{(i)} = (A^{(i)})^\top \lambda^{(i+1)}, i > 0 \tag{258}$$

$$(\eta_+ - \eta_-) = (A^{(0)})^\top \lambda^{(1)} \tag{259}$$

$$\lambda^{(i)} = \frac{u^{(i)}}{u^{(i)} - l^{(i)}}\lambda_-^{(i)} - \lambda_+^{(i)}, i < I \tag{260}$$

$$\lambda^{(I)} = e_j, \tag{261}$$

$$\tag{262}$$

where e_j has a single 1 in entry j and zero in all other entries. We can simplify this by exploiting complementary slackness: we must have $\lambda_+ = \mu_+ = 0$ for all coordinates where $(A^{(i)})^\top \lambda^{(i+1)} > 0$, and similarly $\lambda_- = 0$ for all coordinates where $(A^{(i)})^\top \lambda^{(i+1)} < 0$. Letting $(\cdot)_+$, $(\cdot)_-$ denote the positive and negative parts of a vector,

we then have

$$\lambda_-^{(i)} = ((A^{(i)})^\top \lambda^{(i+1)})_+, \tag{263}$$

$$\lambda_+^{(i)} + \mu_+^{(i)} = ((A^{(i)})^\top \lambda^{(i+1)})_-. \tag{264}$$

$$\tag{265}$$

Since μ_+ is an arbitrary non-negative vector, the latter constraint can be re-expressed as $\lambda_+^{(i)} = \alpha^{(i)}((A^{(i)})^\top \lambda^{(i+1)})_-$ for some vector α whose entries lie in $[0, 1]$. Thus simplifying by substituting in α to remove μ_+, and also merging η_- and η_+ together, we obtain

$$\underset{\lambda,\eta,\lambda_+\geq 0,\lambda_-\geq 0}{\text{minimize}} \ \sum_{i=1}^{I-1} \langle \lambda_-^{(i)}, \frac{u^{(i)} l^{(i)}}{u^{(i)} - l^{(i)}} \rangle + \epsilon \|\eta\|_1 \tag{266}$$

$$\text{subject to } \lambda_-^{(i)} = \max((A^{(i)})^\top \lambda^{(i+1)}, 0), i > 0 \tag{267}$$

$$\lambda_+^{(i)} = \alpha^{(i)} \max(-(A^{(i)})^\top \lambda^{(i+1)}, 0), i > 0 \tag{268}$$

$$\eta = (A^{(0)})^\top \lambda^{(1)}, \tag{269}$$

$$\lambda^{(i)} = \frac{u^{(i)}}{u^{(i)} - l^{(i)}} \lambda_-^{(i)} - \lambda_+^{(i)}, i < I \tag{270}$$

$$\lambda^{(I)} = e_j, \tag{271}$$

$$\tag{272}$$

For fixed $\alpha^{(i)}$, this further simplifies to a simple backpropagation:

$$\underset{\lambda, 0\leq\alpha\leq 1}{\text{minimize}} \ \sum_{i=0}^{I-1} \langle \max((A^{(i)})^\top \lambda^{(i+1)}, 0), \frac{u^{(i)} l^{(i)}}{u^{(i)} - l^{(i)}} \rangle + \epsilon \|(A^{(0)})^\top \lambda^{(1)}\|_1 \tag{273}$$

$$\text{subject to } \lambda^{(i)} = \frac{u^{(i)}}{u^{(i)} - l^{(i)}} \max((A^{(i)})^\top \lambda^{(i+1)}, 0) - \alpha^{(i)} \max(-(A^{(i)})^\top \lambda^{(i+1)}, 0), i < I \tag{274}$$

$$\lambda^{(I)} = e_j. \tag{275}$$

Finally, we can choose the feasible value $\alpha^{(i)} = \frac{u^{(i)}}{u^{(i)} - l^{(i)}}$, in which case λ is determined and evolves according to the particularly simple backpropagation dynamics $\lambda^{(i)} = \frac{u^{(i)}}{u^{(i)} - l^{(i)}} (A^{(i)})^\top \lambda^{(i+1)}$.

We conveniently no longer need to solve a linear program for each node in the network, but can instead obtain the upper and lower bounds (as well as the final bound) just via backpropagation. This still necessitates many backpropagation operations, and the bound could end up being fairly loose. Indeed, for an arbitrarily chosen network this bound will likely be very loose. However, optimizing this bound will tend to find networks where the bound is fairly tight. See Wong and Kolter (2017) for futher details. Also, Dvijotham et al. (2018) train a neural network to *predict* a good choice of the dual variables, rather than guessing them by hand.

Semidefinite programming relaxation. Recall that our goal in the LP relaxation was to handle the constraint $x = \max(z, 0)$. We did this by computing a convex outer bound to this set assuming an upper and lower bound on z. We can instead express $x = \max(z, 0)$ as a system of polynomial constraints and apply an SDP relaxation. Specifically:

$$x = \max(z, 0) \iff x \geq z, \quad x \geq 0, \quad x(x - z) = 0. \tag{276}$$

If we add additional variables X, Y, and Z (meant to represent x^2, xz, and z^2), and let $M = \begin{bmatrix} 1 & x & z \\ x & X & Y \\ z & Y & Z \end{bmatrix}$, we can write these constraints as

$$M \succeq 0, \text{rank}(M) = 1, x \geq z, x \geq 0, X = Y. \tag{277}$$

Dropping the rank constraint and substituting $Y = X$, we obtain the SDP relaxation

$$\begin{bmatrix} 1 & x & z \\ x & X & X \\ z & X & Z \end{bmatrix} \succeq 0, x \geq z, x \geq 0. \tag{278}$$

Applying this to all of the variables in the original optimization problem, we obtain:

$$\underset{x,z,X,Z}{\text{maximize}}\ z_2^{(L)} - z_1^{(L)} \tag{279}$$

$$\text{subject to } z^{(i+1)} = A^{(i)}x^{(i)} \qquad \text{for } i = 0, \ldots, L-1 \tag{280}$$

$$Z^{(i+1)} = A^{(i)}X^{(i)}(A^{(i)})^\top \qquad \text{for } i = 0, \ldots, L-1 \tag{281}$$

$$x^{(i)} \geq z^{(i)}, x^{(i)} \geq 0 \qquad \text{for } i = 0, \ldots, L-1 \tag{282}$$

$$\begin{bmatrix} 1 & (x^{(i)})^\top (z^{(i)})^\top & \\ x^{(i)} & X^{(i)} & X^{(i)} \\ z^{(i)} & X^{(i)} & Z^{(i)} \end{bmatrix} \succeq 0, \tag{283}$$

$$|x_j^{(0)} - x_j| \leq \epsilon \qquad \text{for } j = 1, \ldots, d. \tag{284}$$

For efficiency we would want to collapse x and z (and X and Z) into a single variable by exploiting the equality constraints, but we ignore that for simplicity. There is also one problem with the above optimization, which is that the absolute value constraint on $x^{(0}$ does little to control $X^{(0)}$. A better constraint (which is generally necessary to obtain non-vacuous bounds) is $X_{jj}^{(0)} - 2x_j^{(0)}x_j + x_j^2 \leq \epsilon^2$, which is obtained by squaring the absolute value constraint and applying the SDP relaxation (a good exercise is to show that this implies the original absolute value constraint).

Finally, this relaxation generally performs better if it also makes use of upper and lower bounds (obtained, e.g., via the same method as the LP). If we know that $u \leq x \leq l$, we can incorporate this as the quadratic constraint $(x - l)(x - u) \leq 0$, which becomes $X - (l + u)x + ul \leq 0$.

With these combined improvements, the SDP relaxation is sometimes powerful enough to give good certified bounds even for networks that were not trained against the bound; see Raghunathan et al. (2018) for details.

Comparison of approaches. The LP relaxation is generally faster while the SDP relaxation is tighter. However, an important point in this setting is that *speed can often be turned into accuracy.* This is because faster solve times allow us to train larger neural networks, which have more flexibility to find parameter configurations for which the verifier works well. The LP currently so far has been more successful, although both methods would likely improve with further development. In addition, other verification strategies such as interval bound propagation have also been used to obtain certified bounds. However, the current limit of any of these approaches seems to be medium-sized datasets such as CIFAR-10; perhaps in the future we will figure out how to scale these approaches to ImageNet.

[Lecture 19]

5.4 Randomized Smoothing

We next discuss a simpler, almost trivial approach to obtaining verified bounds, that nevertheless works very well in practice (it currently has the best certified bounds for ℓ_2 perturbations, and is efficient enough to scale to the ImageNet dataset).

The basic idea is as follows: suppose that we have some classifier $f_\theta : \mathbb{R}^d \to [0, 1]^k$, which maps an input $x \in \mathbb{R}^d$ to a k-dimensional vector of class probabilities (so actually the range is $\Delta_k \subset [0, 1]^k$). We can define a *smoothed classifier* $\bar{f}_\theta$ as

$$\bar{f}_\theta(x) = \mathbb{E}_{\delta \sim \pi}[f_\theta(x + \delta)]. \tag{285}$$

In other words, $\bar{f}_\theta$ applies f_θ to some randomly perturbed point $x + \delta$ that is close to x. Observe that we can approximate $\bar{f}_\theta$ well by sampling repeatedly from δ.

Let π_x be the distribution of $x + \delta$. The key bound underlying randomized smoothing lets us control the change in $\bar{f}_\theta$ in terms of a certain modulus of continuity:

Proposition 5.2. *Suppose that f_θ maps into $[0,1]^k$. Then for any x, x', we have*

$$\|\bar{f}_\theta(x) - \bar{f}_\theta(x')\|_\infty \leq \mathsf{TV}(\pi_x, \pi_{x'}). \tag{286}$$

In particular, if $d(x, x') \leq \epsilon$, then $\|\bar{f}_\theta(x) - \bar{f}_\theta(x')\|_\infty \leq \max\{\mathsf{TV}(\pi_x, \pi_{x'}) \mid d(x, x') \leq \epsilon\}$.

This says that $\bar{f}$ is stable under perturbations as long as the family of distributions π_x has bounded modulus of continuity of TV with respect to d (note this is the *opposite* direction of the modulus that we considered before).

The way to apply Proposition 5.2 is to somehow obtain a model such that the probability assigned to the correct class under $\bar{f}_\theta$ is at least τ larger than the probability of any incorrect class. Then as long as $\mathsf{TV}(\pi_x, \pi_{x'}) < \tau$ whenever $d(x, x') \leq \epsilon$, we know that no perturbation can change the arg max prediction of $\bar{f}_\theta$. In the remainder of this section we will discuss how to choose π, and how to train $\bar{f}_\theta$.

Choosing the smoothing distribution π. We will restrict ourselves to the special case $d(x, x') = \|x - x'\|_2$, i.e. ℓ_2 perturbations. In this case we will take $\pi = \mathcal{N}(0, \sigma^2 I)$ for some σ. Then the modulus becomes

$$\max\{\mathsf{TV}(\mathcal{N}(0, \sigma^2 I), \mathcal{N}(\delta, \sigma^2 I)) \mid \|\delta\|_2 \leq \epsilon\} = \Phi(\epsilon/2\sigma) - \Phi(-\epsilon/2\sigma), \tag{287}$$

where Φ is the normal CDF. When ϵ/σ is small, the right-hand-side is $\Theta(\epsilon/\sigma)$, so we are automatically resistant to perturbations that are small in ℓ_2-norm compared to σ.

Observe that the *per-coordinate* noise we apply is comparable in magnitude to the *overall* norm of the perturbation. Thus in d dimensions, we need to apply noise that is $\sqrt{d}$ times larger than the adversarial perturbation that we seek robustness against. This matches the observation on random vs. adversarial noise for linear models from the previous section. Indeed, the above analysis is essentially tight for linear models (up to constants, and assuming ϵ/σ is small).

Training the model. Recalling that f_θ and $\bar{f}_\theta$ both output probability distributions over y, a natural training objective would be to minimize

$$\mathbb{E}_{(x,y)\sim p}[-\log(\bar{f}_\theta(x)_y)] = \mathbb{E}_{(x,y)\sim p}[-\log(\mathbb{E}_\delta[f_\theta(x+\delta)_y])], \tag{288}$$

i.e. the negative log probability that $\bar{f}_\theta(x)$ assigns to the true label y. However, the derivative of this quantity is inconvenient to work with:

$$\nabla_\theta[\log(\mathbb{E}_\delta[f_\theta(x+\delta)_y]] = \mathbb{E}_\delta[\nabla_\theta[f_\theta(x+\delta)_y]]/\bar{f}_\theta(x)_y \tag{289}$$

$$= \mathbb{E}_\delta[\frac{f_\theta(x+\delta)_y}{\bar{f}_\theta(x)_y} \nabla_\theta \log f_\theta(x+\delta)_y]. \tag{290}$$

In particular, the importance weight $\frac{f_\theta(x+\delta)_y}{\bar{f}_\theta(x)_y}$ could have high variation and so require many samples to obtain a good estimate. An alternative is to instead move the log inside the expectation and minimize

$$\mathbb{E}_{(x,y)\sim p}\mathbb{E}_\delta[-\log(f_\theta(x+\delta)_y)]. \tag{291}$$

Then we can compute stochastic gradients of the objective by sampling (x, y), sampling δ, and taking the gradient of $-\log f_\theta(x+\delta)_y$, which can generally be computed straightforwardly (e.g. via backpropagation in the case of neural networks).

[Lecture 20]

6 Domain Adaptation under Covariate Shift

We now shift focus again, to a type of perturbation called *covariate shift*. We work in a classification or regression setting where we wish to predict y from x, and make the assumption that $\tilde{p}(y \mid x)$ and $p^*(y \mid x)$ are the same (the labeling function doesn't change between train and test):

Assumption 6.1 (Covariate Shift)**.** *For a train distribution $\tilde{p}$ and test distribution p^*, we assume that $\tilde{p}(y \mid x) = p^*(y \mid x)$ for all x.*

Thus the only thing that changes between train and test is the distribution of the covariates x (hence the name covariate shift). We furthermore assume that we observe labeled samples $(x_1, y_1), \ldots, (x_n, y_n) \sim \tilde{p}$, together with *unlabeled* samples $\bar{x}_1, \ldots, \bar{x}_m \sim p^*$. In the language of our previous setting, we could say that $D(\tilde{p}, p^*) = \|\tilde{p}(y \mid x) - p^*(y \mid x)\|_\infty$, $\epsilon = 0$, and $\mathcal{G} = \{p \mid p(x) = p_0(x)\}$ for some distribution p_0 (obtained via the unlabeled samples from p^*).

Beyond covariate shift, we will need to make some additional assumption, since if $\tilde{p}(x)$ and $p^*(x)$ have disjoint supports then the assumption that $\tilde{p}(y \mid x) = p^*(y \mid x)$ is meaningless. We will explore two different assumptions:

1. Either we assume the $\tilde{p}(x)$ and $p^*(x)$ are known and not too different from each other, or
2. We assume that the model family is realizable: $p^*(y \mid x) = p_\theta(y \mid x)$ for some θ.

This will lead to two different techniques: importance weighting and uncertainty estimation. We will also see how to construct a "doubly robust" estimator that works as long as at least one of the assumptions holds.

6.1 Importance weighting

First assume that $\tilde{p}(x)$ and $p^*(x)$ are known. (We can generally at least attempt to estimate them from unlabeled data, although if our model family is misspecified then our estimates might be poor.)

In a traditional setting, to minimize the loss on $\tilde{p}(x)$ we would minimize

$$\mathbb{E}_{(x,y)\sim\tilde{p}}[\ell(\theta; x, y)], \tag{292}$$

where ℓ is the loss function for either classification or regression. We can approximate this via the samples from $\tilde{p}$ as

$$\frac{1}{n}\sum_{i=1}^{n} \ell(\theta; x_i, y_i). \tag{293}$$

To handle covariate shift we would like to instead minimize the expectation over p^*, but unfortunately we can't do this because we don't have any outputs y drawn from p^*. The key insight that lets us get around this is the following identity:

$$\mathbb{E}_{(x,y)\sim p^*}\Big[\ell(\theta; x, y)\Big] = \mathbb{E}_{(x,y)\sim\tilde{p}}\Big[\frac{p^*(x)}{\tilde{p}(x)}\ell(\theta; x, y)\Big]. \tag{294}$$

Taking this identity as given for the moment, we can then approximate the expectation over p^* via *samples from $\tilde{p}$* as follows:

$$\frac{1}{n}\sum_{i=1}^{n} \frac{p^*(x_i)}{\tilde{p}(x_i)}\ell(\theta; x_i, y_i). \tag{295}$$

This quantity is called the *propensity-weighted training loss*[4], because each training sample is weighted by how much more it looks like a sample from p^* than from $\tilde{p}$.

[4] Also sometimes called the importance-weighted loss.

To prove the identity, we make use of the covariate shift assumption:

$$\mathbb{E}_{(x,y)\sim p^*}[\ell(\theta;x,y)] = \mathbb{E}_{(x,y)\sim\tilde{p}}\Big[\frac{p^*(x,y)}{\tilde{p}(x,y)}\ell(\theta;x,y)\Big] \tag{296}$$

$$= \mathbb{E}_{(x,y)\sim\tilde{p}}\Big[\frac{p^*(x)}{\tilde{p}(x)}\frac{p^*(y\mid x)}{\tilde{p}(y\mid x)}\ell(\theta;x,y)\Big] \tag{297}$$

$$= \mathbb{E}_{(x,y)\sim\tilde{p}}\Big[\frac{p^*(x)}{\tilde{p}(x)}\ell(\theta;x,y)\Big], \tag{298}$$

where the final equality is by the covariate shift assumption.

Variance of the estimator. Even if $\frac{p^*(x)}{\tilde{p}(x)}$ can be computed, the importance weighted estimator could have high variance. This is because the weights $\frac{p^*(x_i)}{\tilde{p}(x_i)}$ could be large or potentially infinite.

For convenience assume that $\ell(\theta;x,y) \le B$ for all θ, x, y. We can compute (or rather, bound) the variance as follows:

$$\mathsf{Var}_{\tilde{p}}[\frac{p^*(x)}{\tilde{p}(x)}\ell(\theta;x,y)] = \mathbb{E}_{\tilde{p}}[(p^*(x)/\tilde{p}(x))^2\ell(\theta;x,y)^2] - \mathbb{E}_{p^*}[\ell(\theta;x,y)]^2 \tag{299}$$

$$\le \mathbb{E}_{\tilde{p}}[(p^*(x)/\tilde{p}(x))^2]B^2 \tag{300}$$

$$= (D_{\chi^2}(\tilde{p}\|p^*)+1)B^2, \tag{301}$$

where D_{χ^2} is the χ^2-divergence:

$$D_{\chi^2}(\tilde{p}\|p^*) \stackrel{\text{def}}{=} \int \frac{(p^*(x)-\tilde{p}(x))^2}{\tilde{p}(x)}dx = \int \frac{p^*(x)^2}{\tilde{p}(x)}dx - 1. \tag{302}$$

The variance of the propensity-weighted loss is thus more or less controlled by the χ^2-divergence between source and target. To gain some intuition for how χ^2 behaves, first note that is is always larger than KL divergence (in the reverse direction, though I'm not sure the order of arguments is canonical):

$$D_{\mathrm{kl}}(p^*\|\tilde{p}) = \int p^*(x)\log\frac{p^*(x)}{\tilde{p}(x)}dx \tag{303}$$

$$\le \int p^*(x)\frac{p^*(x)-\tilde{p}(x)}{\tilde{p}(x)}dx \tag{304}$$

$$= \int \frac{p^*(x)^2}{\tilde{p}(x)}dx - 1 = D_{\chi^2}(\tilde{p}\|p^*). \tag{305}$$

Additionally, the χ^2-divergence between two Gaussians is exponential in the difference between their means. To see this, let Z denote the normalization constant of an isotropic Gaussian, and write

$$D_{\chi^2}(\mathcal{N}(\mu,I),\mathcal{N}(\mu',I)) = -1 + \frac{1}{Z}\int \exp(\frac{1}{2}\|x-\mu'\|_2^2 - \|x-\mu\|_2^2)dx \tag{306}$$

$$= -1 + \frac{1}{Z}\int \exp(\frac{1}{2}(-\|x\|_2^2 - (2\mu'-4\mu)^\top x + \|\mu'\|_2^2 - 2\|\mu\|_2^2)) \tag{307}$$

$$= -1 + \frac{1}{Z}\int \exp(\frac{1}{2}(-\|x+(\mu'-2\mu)\|_2^2 + \|\mu'-2\mu\|_2^2 + \|\mu'\|_2^2 - 2\|\mu\|_2^2)) \tag{308}$$

$$= -1 + \exp(\|\mu'\|_2^2 + \|\mu\|_2^2 - 2\mu^\top\mu') = -1 + \exp(\|\mu-\mu'\|_2^2). \tag{309}$$

This is bad news for propensity weighting, since the weights blow up exponentially as the distributions move apart.

Connection to causal inference. Propensity weighting can also be used in the context of causal inference. Here we have a patient with covariates X, with treatment condition T (usually $T \in \{0, 1\}$), and outcome Y. Our goal is to estimate the treatment effect, which, roughly speaking, is $\mathbb{E}[Y \mid T = 1] - \mathbb{E}[Y \mid T = 0]$ (this is wrong as stated and will be remedied below). We will see below how to do this by letting p_0^* and p_1^* be the distributions where $T = 0$ and $T = 1$, respectively. However, first we need to set up the problem more carefully.

To set the problem up more carefully, we use the *potential outcomes framework*. In this framework there are actually two variables, $Y(0)$ and $Y(1)$, which are what the outcome *would have been* if we had set $T = 0$ or $T = 1$, respectively. This is potentially different from the distribution of the outcome conditional on T, since there could be factors that correlate T with Y (for instance, if T is smoking and Y is lung cancer, there could be some gene that causes one to both be more likely to smoke and more likely to get lung cancer that accounts for the strong empirical correlation between T and Y; this was an actual objection raised by Fisher!).

Of course, there are plenty of factors that create correlation between T and Y in an observational setting, for instance sicker patients are more likely to be treated aggressively. We are okay with this as long as these factors are observed as part of the covariates X. This leads us to the *unconfoundedness assumption*:

Assumption 6.2 (Unconfoundedness)**.** *The distribution $(X, T, Y(0), Y(1))$ is said to be* unconfounded *if $Y(0), Y(1) \perp\!\!\!\perp T \mid X$. In other words, treatment and outcome should be independent conditional on the covariates X.*

The main challenge in the potential outcomes framework is that we only observe $(X, T, Y(T))$. In other words, we only observe the outcome for the treatment T that was actually applied, which makes it difficult to estimate $\mathbb{E}[Y(1)]$ or $\mathbb{E}[Y(0)]$. We will deal with this by treating estimating $\mathbb{E}[Y(1)]$ as a domain adaptation problem, and using propensity weighting. First note that, by unconfoundedness, we have

$$\begin{aligned}
\mathbb{E}_{\tilde{p}}[Y(1)] &= \mathbb{E}_{X \sim \tilde{p}}[\mathbb{E}_{\tilde{p}}[Y(1) \mid X]] && (310)\\
&= \mathbb{E}_{X \sim \tilde{p}}[\mathbb{E}_{\tilde{p}}[Y(1) \mid X, T = 1]] && (311)\\
&= \mathbb{E}_{p_1^*}[Y(T)], && (312)
\end{aligned}$$

where we define p_1^* such that $p_1^*(x, t, y) = \tilde{p}(x)\mathbb{I}[t = 1]\tilde{p}(y \mid x, t = 1)$; this has the same distribution over x as $\tilde{p}$, but the treatment $t = 1$ is always applied. Since $\tilde{p}(y \mid x, t) = p^*(y \mid x, t)$ almost surely, the covariate shift assumption holds. We can thus estimate the expectation under p_1^* via propensity weighting:

$$\begin{aligned}
\mathbb{E}_{p_1^*}[Y(T)] &= \mathbb{E}_{\tilde{p}}\Big[\frac{p_1^*(X, T)}{\tilde{p}(X, T)} Y(T)\Big] && (313)\\
&= \mathbb{E}_{\tilde{p}}\Big[\frac{p_1^*(T \mid X)}{\tilde{p}(T \mid X)} Y(T)\Big] && (314)\\
&= \mathbb{E}_{\tilde{p}}\Big[\frac{\mathbb{I}[T = 1]}{\tilde{p}(T \mid X)} Y(T)\Big]. && (315)
\end{aligned}$$

A similar calculation holds for computing $\mathbb{E}_{\tilde{p}}[Y(0)]$, for the distribution $p_0^*(x, t, y) = \tilde{p}(x)\mathbb{I}[t = 0]\tilde{p}(y \mid x, t = 0)$. Together, we have that

$$\mathbb{E}_{\tilde{p}}[Y(1) - Y(0)] = \mathbb{E}_{\tilde{p}}\Big[\Big(\frac{\mathbb{I}[T = 1]}{\tilde{p}(T \mid X)} - \frac{\mathbb{I}[T = 0]}{\tilde{p}(T \mid X)}\Big) Y(T)\Big]. \tag{316}$$

Since the right-hand-side is in terms of $Y(T)$, it only involves observable quantities, and can be estimated from samples as long as $\tilde{p}(T \mid X)$ is known. This estimator is called *inverse propensity weighting* because it involves dividing by the propensity weights $\tilde{p}(T \mid X)$.

In the next section, we will explore an improvement on inverse propensity weighting called a *doubly-robust estimator*.

[Lecture 21]

6.2 Doubly-Robust Estimators

Recall that in the previous section we defined the inverse propensity weighted estimator

$$\mathbb{E}_{\tilde{p}}[Y(1) - Y(0)] = \mathbb{E}_{\tilde{p}}\Big[\Big(\frac{\mathbb{I}[T=1]}{\tilde{p}(T \mid X)} - \frac{\mathbb{I}[T=0]}{\tilde{p}(T \mid X)}\Big)Y(T)\Big]. \tag{317}$$

To actually estimate the left-hand-side, we take the empirical average over n samples.

There are a couple of downsides of this estimator. One is that the variance of this estimator can be large. Specifically, we can compute it as

$$\frac{1}{n}\Big(\mathbb{E}_{\tilde{p}}\Big[\frac{1}{\tilde{p}(T=1 \mid X)}Y(1)^2 + \frac{1}{\tilde{p}(T=0 \mid X)}Y(0)^2\Big] - \mathbb{E}_{\tilde{p}}[Y(1) - Y(0)]^2\Big). \tag{318}$$

If the propensity weights are near zero then the variance explodes (similarly to the issue with χ^2-divergence that we saw earlier).

Another issue is that estimating the propensity weights themselves is non-trivial, and if we use the wrong propensity weights, then the estimate could be arbitrarily wrong.

We will explore an idea that partially mitigates both issues; it reduces the variance when the propensity weights are correct (although doesn't avoid the exploding variance issue), and in some cases it produces a correct estimate even if the propensity weights are wrong.

The basic idea is as follows: suppose that we have some prediction $\bar{Y}(1, X)$, $\bar{Y}(0, X)$ of what will happen under $T = 1$, $T = 0$ conditioned on X. Since these predictions only require knowing X and not T, an alternate estimate of the treatment effect can be obtained by adding and subtracting $\bar{Y}$:

$$\mathbb{E}_{\tilde{p}}[Y(1) - Y(0)] = \mathbb{E}_{\tilde{p}}[\bar{Y}(1, X) - \bar{Y}(0, X)] + \mathbb{E}_{\tilde{p}}[(Y(1) - \bar{Y}(1, X)) - (Y(0) - \bar{Y}(0, X))] \tag{319}$$

$$= \mathbb{E}_{\tilde{p}}[\bar{Y}(1, X) - \bar{Y}(0, X)] + \mathbb{E}_{\tilde{p}}\Big[\Big(\frac{\mathbb{I}[T=1]}{\tilde{p}(T \mid X)} - \frac{\mathbb{I}[T=0]}{\tilde{p}(T \mid X)}\Big)(Y(T) - \bar{Y}(T, X))\Big]. \tag{320}$$

In other words, we first use our prediction $\bar{Y}$ to form a guess of the average treatment effect, then use inverse propensity weighting to correct the guess so as to obtain an unbiased estimate. This can yield substantial improvements when $Y(T) - \bar{Y}(T, X)$ is much smaller in magnitude than $Y(T)$. For instance, a patient's cholesterol after taking a cholesterol-reducing drug is still highly-correlated with their initial cholesterol, so in that case we can take $\bar{Y}(T, X)$ to be the pre-treatment cholesterol level. Even though this is independent of T it can substantially reduce the variance of the estimate! (We will formally bound the variance below.)

Bias of the estimate. Call the prediction $\bar{Y}$ unbiased if $\mathbb{E}[Y \mid X, T] = \bar{Y}(T, X)$. The first key property of (320) is that it is unbiased as long as *either* $\bar{Y}$ is unbiased, or the propensity weights are correct. Indeed, if the prediction is unbiased then the first term is the average treatment effect while the second term is zero. Conversely, if the propensity weights are correct then the second term exactly estimates the difference between the predicted and true treatment effect. Correspondingly, (320) is called a *doubly-robust estimator*.

We can actually say more about the bias. Suppose that instead of the true propensity weights, we have an incorrect guess $\hat{p}(t \mid x)$. Then the bias of the estimate is the difference between $\mathbb{E}_{\tilde{p}}[Y(1) - Y(0)]$ and (320), which is

$$\mathbb{E}_{\tilde{p}}[Y(1) - Y(0)] - \mathbb{E}[\bar{Y}(1, X) - \bar{Y}(0, X)] - \mathbb{E}_{\tilde{p}}\Big[\Big(\frac{\mathbb{I}[T=1]}{\hat{p}(T \mid X)} - \frac{\mathbb{I}[T=0]}{\hat{p}(T \mid X)}\Big)(Y(T) - \bar{Y}(T, X))\Big] \tag{321}$$

$$= \mathbb{E}_{\tilde{p}}\Big[(Y(1) - \bar{Y}(1, X))\Big(1 - \frac{\mathbb{I}[T=1]}{\hat{p}(t=1 \mid X)}\Big) + (Y(0) - \bar{Y}(0, X))\Big(1 - \frac{\mathbb{I}[T=0]}{\hat{p}(t=0 \mid X)}\Big)\Big]. \tag{322}$$

Focusing on the first term, and using the independence of T and Y conditioned on X, we have

$$\mathbb{E}_{\tilde{p}}\Big[(Y(1) - \bar{Y}(1, X))\Big(1 - \frac{\mathbb{I}[T=1]}{\hat{p}(t=1 \mid X)}\Big)\Big] \tag{323}$$

$$= \mathbb{E}_{\tilde{p}}\Big[(\mathbb{E}[Y(1) \mid X] - \bar{Y}(1, X))\Big(1 - \frac{\tilde{p}(t=1 \mid X)}{\hat{p}(t=1 \mid X)}\Big) \mid X\Big] \tag{324}$$

$$\leq \mathbb{E}_{\tilde{p}}[(\mathbb{E}[Y(1) \mid X] - \bar{Y}(1, X))^2]^{1/2}\mathbb{E}_{\tilde{p}}\Big[\Big(1 - \frac{\tilde{p}(t=1 \mid X)}{\hat{p}(t=1 \mid X)}\Big)^2\Big]^{1/2}, \tag{325}$$

meaning that the bias of the estimator is the *product* of the biases of $\bar{Y}$ and $\hat{p}$ (measured as the expected squared errors in (325)).

Variance of the estimate. We can obtain a somewhat similar relation for the variance. Usually the variance of $\bar{Y}(1,X) - \bar{Y}(0,X)$ is small compared to the propensity-weighted term, so again focusing on the $T=1$ case we have

$$\mathsf{Var}\Big[(Y(1) - \bar{Y}(1,X))\frac{\mathbb{I}[T=1]}{\hat{p}(t=1 \mid X)}\Big] \leq \mathbb{E}\Big[\mathbb{E}_Y[(Y(1) - \bar{Y}(1,X))^2 \mid X]\frac{\tilde{p}(t=1 \mid X)}{\hat{p}(t=1 \mid X)^2}\Big]. \tag{326}$$

The variance is substantially reduced when $Y(1)$ is close to $\bar{Y}(1,X)$. We cannot always hope for this, e.g. if Y has a large amount of intrinsic variance even conditioned on X. But in many cases even trivial $\bar{Y}$ can predict most of the variance in Y—for instance, for any chronic disease the patient's post-treatment status is well-predicted by their pre-treatment status. And the value of a stock tomorrow is well-predicted by its value today.

Semi-parametric estimation. It may seem difficult to estimate $\bar{Y}(\cdot,X)$ and $\tilde{p}(t=1 \mid X)$, since any parametric model could be mis-specified and lead to biased estimates. One idea is to estimate these both non-parametrically, and then apply the doubly-robust estimator above. This is an instance of *semi-parametric estimation*, because while we estimate $\bar{Y}$ and $\tilde{p}(t \mid X)$ non-parametrically, the doubly-robust estimator itself is parametric (i.e a simple sample estimate of the mean), and in some cases we obtain non-parametric rates. This is explored in detail in Nie and Wager (2017) for estimating conditional average treatment effects; we describe the basic idea here. Since the squared error in an estimate is $\mathsf{Bias}^2 + \mathsf{Variance}/n$, the bias will dominate the error in the doubly-robust estimator as long as the variance doesn't explode (of course, the variance can often explode if the propensity weights are too close to 0 or 1, and the following idea won't help in that case).

We saw above that the bias of the doubly-robust estimator is the product of the biases in $\bar{Y}$ and $\hat{p}$, which are both given as expected squared errors between the true and estimated value. In non-parametric estimation, we typically get convergence rates of $\mathcal{O}(n^{-\alpha})$ for some $\alpha < 1/2$ (note that $\alpha = 1/2$ is what we typically get for parametric estimation). The parameter α typically depends on the dimension of the problem and the smoothness of the function class we wish to estimate. Since the doubly-robust bias is the product of the biases, we end up with a bias of $\mathcal{O}(n^{-2\alpha})$ as long as $\bar{Y}$ and $\hat{p}$ each converge at a $n^{-\alpha}$ rate. As long as $\alpha > 1/4$, this yields a parametric rate (the variance term will then asymptotically dominate as it only converges at $1/\sqrt{n}$).

[Lecture 22]

6.3 Partial Specification for Linear Regression

So far we have made no assumptions about the relation between Y and X, and consequently have only been able to handle small shifts in the distribution $p(x)$ (i.e. $D_{\chi^2}(\tilde{p}(x)\|p^*(x))$ must be small). We will next make very strong assumptions about how Y relates to X, and handle much larger shifts. Then we will try to relax those assumptions using an idea called *partial specification*. We will move away from the causal inference setting above, and instead consider linear regression: we wish to predict $Y \in \mathbb{R}$ from $X \in \mathbb{R}^d$ using some linear predictor $\langle \beta, X \rangle$. As before we make the covariate shift assumption that $\tilde{p}(y \mid x) = p^*(y \mid x)$, but $\tilde{p}(x)$ and $p^*(x)$ could differ. Our goal, rather than to construct a good predictor on p^*, is to estimate the error of the ordinary least squares estimator.

Starting point: linear response with Gaussian errors. In the simplest setting, suppose that we completely believe our model:

$$Y = \langle \beta, X \rangle + Z, \text{ where } Z \sim \mathcal{N}(0, \sigma^2 I). \tag{327}$$

We observe samples $(x_1, y_1), \ldots, (x_n, y_n) \sim \tilde{p}$, and samples $\bar{x}_1, \ldots, \bar{x}_m \sim p^*(x)$. Suppose that we estimate β using the ordinary least squares estimator:

$$\hat{\beta} = \arg\min_{\beta} \frac{1}{n} \sum_{i=1}^{n} (y_i - \langle \beta, x_i \rangle)^2 = (\sum_{i=1}^{n} x_i x_i^\top)^{-1} \sum_{i=1}^{n} x_i y_i. \tag{328}$$

Define $\tilde{\Omega} = \frac{1}{n} \sum_{i=1}^{n} x_i x_i^\top$. Then since $y_i = x_i^\top \beta + z_i$, we can further write

$$\hat{\beta} = (\sum_{i=1}^{n} x_i x_i^\top)^{-1} (\sum_{i=1}^{n} x_i x_i^\top \beta + x_i z_i) \tag{329}$$

$$= (n\tilde{\Omega})^{-1} (n\tilde{\Omega}\beta + \sum_{i=1}^{n} x_i z_i) \tag{330}$$

$$= \beta + \frac{1}{n} \tilde{\Omega}^{-1} \sum_{i=1}^{n} x_i z_i. \tag{331}$$

From this we see that, conditional on the x_i, $\hat{\beta} - \beta$ is a zero-mean Gaussian distribution. Its covariance matrix is given by

$$\frac{1}{n^2} \tilde{\Omega}^{-1} \sum_{i=1}^{n} \mathbb{E}[z_i^2 \mid x_i] x_i x_i^\top \tilde{\Omega}^{-1} = \frac{\sigma^2}{n} \tilde{\Omega}^{-1}. \tag{332}$$

Now suppose that we wish to estimate the error on the samples $\bar{x}_{1:m}$. The expected error on sample $\bar{x}_i$ is $\sigma^2 + \langle \hat{\beta} - \beta, \bar{x}_i \rangle^2$. If we let $\Omega^* = \frac{1}{m} \sum_{i=1}^{m} \bar{x}_i \bar{x}_i^\top$, then the overall average expected error (conditional on $x_{1:n}, \bar{x}_{1:m}$) is

$$\sigma^2 + \mathbb{E}_Z[\frac{1}{m} \sum_{i=1}^{m} (\bar{x}_i^\top (\beta - \hat{\beta}))^2] = \sigma^2 + \langle \frac{1}{m} \sum_{i=1}^{m} \bar{x}_i \bar{x}_i^\top, \mathbb{E}_Z[(\beta - \hat{\beta})(\beta - \hat{\beta})^\top] \rangle \tag{333}$$

$$= \sigma^2 + \langle \Omega^*, \frac{\sigma^2}{n} \tilde{\Omega}^{-1} \rangle \tag{334}$$

$$+ \sigma^2 \Big(1 + \frac{1}{n} \langle \Omega^*, \tilde{\Omega}^{-1} \rangle \Big). \tag{335}$$

This shows that the error depends on the divergence between the second moment matrices of $\tilde{p}(x)$ and $p^*(x)$:

- When $\tilde{p}(x) = p^*(x)$, then $\langle \Omega^*, \tilde{\Omega}^{-1} \rangle = \mathrm{tr}(\Omega^* \tilde{\Omega}^{-1}) \approx \mathrm{tr}(I) = d$, so the error decays as $\frac{d}{n}$.
- If $\tilde{\Omega}$ is low-rank and is missing any directions that appear in Ω, then the error is infinite. This makes sense, as we have no way of estimating β along the missing directions, and we need to be able to estimate β in those directions to get good error under p^*. We can get non-infinite bounds if we further assume some norm bound on β^*; e.g. if $\|\beta^*\|_2$ is bounded then the missing directions only contribute some finite error.
- On the other hand, if $\tilde{\Omega}$ is full-rank but Ω^* is low-rank, then we still achieve finite error. For instance, suppose that $\tilde{\Omega} = I$ is the identity, and $\Omega^* = \frac{d}{k} P$ is a projection matrix onto a k-dimensional subspace, scaled to have trace d. Then we get a sample complexity of $\frac{d}{n}$, although if we had observed samples with second moment matrix Ω^* at training time, we would have gotten a better sample complexity of $\frac{k}{n}$.
- In general it is always better for $\tilde{\Omega}$ to be bigger. This is partially an artefact of the noise σ^2 being the same for all X, so we would always rather have X be as far out as possible since it pins down β more effectively. If the noise was proportional to $\|X\|$ (for instance) then the answer would be different.

Of course, this all so far rests on the assumption of Gaussian error. Can we do better?

Calculation from moment assumptions. It turns out that our calculation above relied only on conditional moments of the errors, rather than Gaussianity. We will show this explicitly by doing the calculations more carefully. Re-using steps above, we have that

$$\hat{\beta} - \beta = \frac{1}{n}\tilde{\Omega}^{-1}\sum_{i=1}^{n} x_i z_i. \tag{336}$$

In particular, assuming that the (x_i, y_i) are i.i.d., we have

$$\mathbb{E}[\hat{\beta} - \beta \mid x_1, \ldots, x_n] = \frac{1}{n}\tilde{\Omega}^{-1}\sum_{i=1}^{n} x_i \mathbb{E}[z_i \mid x_i] = \tilde{\Omega}^{-1}\tilde{b}, \tag{337}$$

where $\tilde{b} \stackrel{\text{def}}{=} \frac{1}{n}\sum_{i=1}^{n} x_i \mathbb{E}[z_i \mid x_i]$.

In particular, as long as $\mathbb{E}[Z \mid X] = 0$ for all X, $\hat{\beta}$ is an unbiased estimator for X. In fact, since this only needs to hold on average, as long as $\mathbb{E}[ZX] = 0$ (covariates uncorrelated with noise) then $\mathbb{E}[\hat{\beta} - \beta] = 0$, and $\mathbb{E}[\hat{\beta} - \beta \mid x_{1:n}]$ converges to zero as $n \to \infty$. This yields an insight that is important more generally:

> *Orindary least squares yields an unbiased estimate of β whenever the covariates X and noise Z are uncorrelated.*

This partly explains the success of OLS compared to other alternatives (e.g. penalizing the absolute error or fourth power of the error). While OLS might initially look like the maximum likelihood estimator under Gaussian errors, it yields consistent estimates of β under much weaker assumptions. Minimizing the fourth power of the error requires stronger assumptions for consistency, while minimizing the absolute error would yield a different condition in terms of medians rather than expectations.

Next we turn to the covariance of $\hat{\beta}$. Assuming that (x_i, y_i) are independent like before, we have

$$\mathsf{Cov}[\hat{\beta} \mid x_{1:n}] = \mathsf{Cov}[\frac{1}{n}\tilde{\Omega}^{-1}\sum_{i=1}^{n} x_i z_i \mid x_{1:n}] \tag{338}$$

$$= \frac{1}{n^2}\tilde{\Omega}^{-1}\sum_{i,j=1}^{n} x_i \mathsf{Cov}[z_i, z_j \mid x_i, x_j] x_j^\top \tilde{\Omega}^{-1} \tag{339}$$

$$= \frac{1}{n^2}\tilde{\Omega}^{-1}\sum_{i=1}^{n} x_i \mathsf{Var}[z_i \mid x_i] x_i^\top \tilde{\Omega}^{-1}, \tag{340}$$

where the final line is because z_i, z_j are independent for $i \neq j$. If we define $\tilde{M} = \frac{1}{n}\sum_{i=1}^{n} x_i \mathsf{Var}[z_i \mid x_i] x_i^\top$, then the final term becomes $\frac{1}{n}\tilde{\Omega}^{-1}\tilde{M}\tilde{\Omega}$.

This quantity is bounded under much weaker assumptions than Gaussianity. If we, for instance, merely assume that $\mathsf{Var}[z_i \mid x_i] \leq \sigma^2$ for all i, then we have that $\tilde{M} \preceq \sigma^2\tilde{\Omega}$ and hence $\mathsf{Cov}[\hat{\beta} \mid x_{1:n}] \preceq \frac{\sigma^2}{n}\tilde{\Omega}^{-1}$.

We can put this together to estimate the squared error. Letting $\bar{z}_i$ be the noise for $\bar{x}_i$, the squared error is then $\frac{1}{m}\sum_{j=1}^{m}(\langle \beta - \hat{\beta}, \bar{x}_i\rangle + \bar{z}_i)^2$, and computing the expectation given $x_{1:n}, \bar{x}_{1:m}$ yields

$$\mathbb{E}[\frac{1}{m}\sum_{j=1}^{m}(\langle \beta - \hat{\beta}, \bar{x}_i\rangle + \bar{z}_i)^2 \mid x_{1:n}, \bar{x}_{1:m}] \tag{341}$$

$$= \frac{1}{m}\sum_{i=1}^{m} \bar{x}_i^\top \mathbb{E}[(\beta - \hat{\beta})(\beta - \hat{\beta})^\top \mid x_{1:n}]\bar{x}_i + 2\bar{x}_i^\top \mathbb{E}[\beta - \hat{\beta} \mid x_{1:n}]\mathbb{E}[\bar{z}_i \mid x_i] + \mathbb{E}[\bar{z}_i^2 \mid \bar{x}_i] \tag{342}$$

$$= \left\langle \Omega^*, \tilde{\Omega}^{-1}\left(\frac{1}{n}\tilde{M} + \tilde{b}\tilde{b}^\top\right)\tilde{\Omega}^{-1}\right\rangle + 2\left\langle b^*, \tilde{\Omega}^{-1}\tilde{b}\right\rangle + \frac{1}{m}\sum_{j=1}^{m}\mathbb{E}[\bar{z}_i^2 \mid \bar{x}_i]. \tag{343}$$

To interpret this expression, first assume that the true model is "actually linear", meaning that $\tilde{b} = b^* = 0$. Then the expression reduces to $\frac{1}{n}\langle \Omega^*, \tilde{\Omega}^{-1}\tilde{M}\tilde{\Omega}^{-1}\rangle + \frac{1}{m}\sum_{j=1}^{m}\mathbb{E}[\bar{z}_i^2 \mid x_i]$. The second term is the intrinsic

variance in the data, while the first term is similar to the $\frac{1}{n}\langle \Omega^*, \tilde{\Omega}^{-1} \rangle$ term from before, but accounts for correlation between X and the variation in Z. The $\tilde{M}$ term is also reminiscent of our earlier condition for robust linear regression.

If the model is not actually linear, then we need to decide how to define β (since the optimal β is then no longer independent of the distribution). In that case a natural choice is to let β be the minimizer under the training distribution, in which case $\tilde{b} \to 0$ as $n \to \infty$ and thus the $\langle b^*, \tilde{\Omega}^{-1}\tilde{b} \rangle$ term conveniently becomes asymptotically negligible. The twist is that $\mathbb{E}[\bar{z}_i^2 \mid \bar{x}_i]$ now measures not just the intrinsic variance but also the departure from linearity, and could be quite large if the linear extrapolation away from the training points ends up being poor.

Partial specification. In general, we see that we can actually form good estimates of the mean-squared error on p^* making only certain moment assumptions (e.g. $\tilde{b} = b^* = 0$) rather than needing to assume the Gaussian model is correct. This idea is called *partial specification*, where rather than making assumptions that are stringent enough to specify a parametric family, we make weaker assumptions that are typically insufficient to even yield a likelihood, but show that our estimates are still valid under those weaker assumptions. The weaker the assumptions, the more happy we are. Of course $\tilde{b} = b^* = 0$ is still fairly strong, but much better than Gaussianity. The goal of partial specification aligns with our earlier desire to design estimators for the entire family of resilient distributions, rather than specific parametric classes. We will study other variants of partial specification in the next section.

[Lecture 23]

7 Agnostic Clustering

We next study the idea of partial specification for clustering. Our setting for clustering will be the following:

- There are k unknown distributions $p_1, \dots, p_k$.
- We observe points $x_1, \dots, x_n$, such that a fraction α_j of the points x_i are drawn from p_j.

Generally the α_j are not known but we have a lower bound on $\alpha_{\min} = \min_{j=1}^k \alpha_j$. In clustering we have two goals:

- **Parameter recovery:** We wish to estimate some parameter of the p_j (usually their means).
- **Cluster recovery:** We wish to determine for each point x_i which cluster p_j it was drawn from.

In the simplest setting, we assume that each of the p_j has a known parametric form (for instance, each p_j is a Gaussian with unknown mean and variance). In the *agnostic* setting, we do not assume a parametric form for the p_j but instead only assume e.g. bounded moments. In the *robust* setting, we allow some fraction ϵ of the points to be arbitrary outliers (so $\alpha_1 + \cdots + \alpha_k = 1 - \epsilon$).

Partial specification thus corresponds to the agnostic setting. Clustering is a particularly interesting setting for studying partial specification because some algorithms that work in the simple setting fail completely in the agnostic setting. Below we will first study the simple setting and give an algorithm based on the method of moments, then turn our attention to the agnostic setting. In the agnostic setting, resilience will appear once again as an information-theoretically sufficient condition enabling clustering. Finally, we will turn our attention to efficient algorithms. In many cases the agnostic algorithms will work even in the robust agnostic setting.

7.1 Clustering Mixtures of Gaussians

Here we assume that each $p_j = \mathcal{N}(\mu_j, \Sigma_j)$. Thus we can treat each x_i as being drawn from $p = \sum_{j=1}^k \alpha_j \mathcal{N}(\mu_j, \Sigma_j)$. This is a parametric model with parameters $(\alpha_j, \mu_j, \Sigma_j)$, so (at least in the limit of infinite data) a sufficient condition for exact parameter recovery is for the model to be identifiable, meaning that if $\sum_{j=1}^k \alpha_j \mathcal{N}(\mu_j, \Sigma_j) = \sum_{j=1}^k \alpha'_j \mathcal{N}(\mu'_j, \Sigma'_j)$, then $\alpha_j = \alpha'_j$, $\mu_j = \mu'_j$, and $\Sigma_j = \Sigma'_j$.[5]

[5] We also need to worry about the case where $k \neq k'$, but for simplicity we ignore this.

As stated, the model is never identifiable because we can always permute the $(\alpha_j, \mu_j, \Sigma_j)$ and obtain an identical distribution. What we actually care about is *identifiability up to permutation*: if $p_{\alpha,\mu,\Sigma} = p_{\alpha',\mu',\Sigma'}$ then $\alpha_j = \alpha'_{\sigma(j)}$, $\mu_j = \mu'_{\sigma(j)}$, and $\Sigma_j = \Sigma'_{\sigma(j)}$ for some permutation σ.

We have the following result:

Proposition 7.1. *As long as the orders pairs (μ_j, Σ_j) are all distinct, the parameters $(\alpha_j, \mu_j, \Sigma_j)$ are identifiable up to permutation.*

Proof. This is equivalent to showing that the functions $f_{\mu,\Sigma}(x)$ defining the pdf of a Gaussian are all linearly independent (i.e., there is no non-trivial finite combination that yields the zero function). We will start by showing this in one dimension. So, suppose for the sake of contradiction that

$$\sum_{j=1}^{m} c_j \exp(-(x-\mu_j)^2/2\sigma_j^2)/\sqrt{2\pi\sigma^2} = 0, \tag{344}$$

where the c_j are all non-zero. Then integrating (344) against the function $\exp(\lambda x)$ and using the formula for the moment generating function of a Gaussian, we obtain

$$\sum_{j=1}^{m} c_j \exp(\frac{1}{2}\sigma_j^2\lambda^2 + \mu_j\lambda) = 0. \tag{345}$$

Let $\sigma_{\max} = \max_{j=1}^{m} \sigma_j$, then dividing the above equation by $\exp(\frac{1}{2}\sigma_{\max}^2\lambda^2)$ and taking $\lambda \to \infty$, we see that only those j such that $\sigma_j = \sigma_{\max}$ affect the limit. If S is the set of such indices j, we obtain

$$\sum_{j\in S} c_j \exp(\mu_j\lambda) = 0, \tag{346}$$

i.e. there is a linear relation between the functions $g_{\mu_j}(\lambda) = \exp(\mu_j\lambda)$. But this is impossible, because as long as the μ_j are distinct, the largest μ_j will always dominate the limit of the linear relation as $\lambda \to \infty$, and so we must have $c_j = 0$ for that j, a contradiction.

It remains to extend to the n-dimensional case. Suppose there was a linear relation among the PDFs of n-dimensional Gaussians with distinct parameters. Then if we project to a random 1-dimensional subspace, the corresponding marginals (which are linear functions of the n-dimensional PDFs) are also each Gaussian, and have distinct parameters with probability 1. This is again a contradiction since we already know that distinct 1-dimensional Gaussians cannot satisfy any non-trivial linear relation. □

Proposition 7.1 shows that we can recover the parameters exactly in the limit of infinite data, but it doesn't say anything about finite-sample rates. However, asymptotically, as long as the log-likelihood function is locally quadratic around the true parameters, we can use tools from asymptotic statistics to show that we approach the true parameters at a $1/\sqrt{n}$ rate.

Recovery from moments. Proposition 7.1 also leaves open the question of efficient computation. In practice we would probably use k-means or EM, but another algorithm is based on the *method of moments*. It has the virtue of being provably efficient, but is highly brittle to mis-specification.

The idea is that the first, second, and third moments give a system of equations that can be solved for the parameters (α, μ, Σ): letting $p = \sum_j \alpha_j \mathcal{N}(\mu_j, \Sigma_j)$, we have

$$\mathbb{E}_p[X] = \sum_{j=1}^{k} \alpha_j \mu_j, \tag{347}$$

$$\mathbb{E}_p[X \otimes X] = \sum_{j=1}^{k} \alpha_j (\mu_j \mu_j^\top + \Sigma_j), \tag{348}$$

$$\mathbb{E}_p[X \otimes X \otimes X] = \sum_{j=1}^{k} \alpha_j (\mu_j^{\otimes 3} + 3\,\mathrm{Sym}(\mu_j \otimes \Sigma_j)), \tag{349}$$

where $\mathrm{Sym}(X)_{i_1 i_2 i_3} = \frac{1}{6}(X_{i_1 i_2 i_3} + X_{i_1 i_3 i_2} + X_{i_2 i_1 i_3} + X_{i_2 i_3 i_1} + X_{i_3 i_1 i_2} + X_{i_3 i_2 i_1})$.

In d dimensions, this yields $d + \binom{d+1}{2} + \binom{d+2}{3} \approx d^3/6$ equations and $k(1 + d + \binom{d+1}{2}) \approx kd^2/2$ unknowns. Thus as long as $d > 3k$ we might hope that these equations have a unique (up to permutation) solution for (α, μ, Σ). As an even more special case, if we assume that the covariance matrices are all diagonal, then we only have approximately $2kd$ unknowns, and the equations have a solution whenever the μ_j are linearly independent. We can moreover find this solution via an efficient algorithm called the *tensor power method*, which is a generalization of the power method for matrices, and the rate of convergence is polynomial in k, d, and the condition number of certain matrices (and decays as $1/\sqrt{n}$).

However, this method is very brittle—it relies on exact algebraic moment relations of Gaussians, so even small departures from the assumptions (like moving from Gaussian to sub-Gaussian) will likely break the algorithm. This is one nice thing about the agnostic clustering setting—it explicitly reveals the brittleness of algorithms like the one above, and (as we shall see) shows why other algorithms such as k-means are likely to perform better in practice.

Cluster recovery. An important point is that even in this favorable setting, exact cluster recovery is impossible. This is because even if the Gaussians are well-separated, there is some small probability that a sample ends up being near the center of a different Gaussian.

To measure this quantitatively, assume for simplicity that $\Sigma_j = \sigma^2 I$ for all j (all Gaussians are isotropic with the same variance), and suppose also that the μ_j are known exactly and that we assign each point x to the cluster that minimizes $\|x - \mu_j\|_2$.[6] Then the error in cluster recovery is exactly the probability that a sample from μ_j ends up closer to some other sample $\mu_{j'}$, which is

$$\sum_{j=1}^{k} \alpha_j \mathbb{P}_{x \sim \mathcal{N}(\mu_j, \sigma^2 I)}[\|x - \mu_j\|_2 > \|x - \mu_{j'}\|_2 \text{ for some } j' \neq j] \leq \sum_{j=1}^{k} \alpha_j \sum_{j' \neq j} \Phi(\|\mu_j - \mu_{j'}\|/\sigma) \tag{350}$$

$$\leq k\Phi(\Delta/\sigma), \tag{351}$$

where $\Delta = \min_{j' \neq j} \|\mu_j - \mu_{j'}\|_2$ and Φ is the normal CDF. As long as $\Delta \gg \sqrt{\log(k/\epsilon)}$, the cluster error will be at most ϵ. Note that the cluster error depends on a *separation condition* stipulating that the cluster centers are all sufficiently far apart. Moreover, we need greater separation if there are more total clusters (albeit at a slowly-growing rate in the Gaussian case).

[Lecture 24]

7.2 Clustering Under Resilience

The mixture of Gaussians case is unsatisfying because data are unlikely to actually be Gaussian mixtures in practice, yet common algorithms like k-means still do a good job at clustering data. We therefore move to the agnostic setting, and show that we only need the distributions to be *resilient* in order to cluster successfully.

We will start by proving an even stronger result—that if a set of points contains a (ρ, α)-resilient subset S of size αn, then it is possible to output an estimate $\hat{\mu}$ that is close to the true mean μ of S, regardless of the other $(1 - \alpha)n$ points. As stated, this is impossible, since there could be $\mathcal{O}(1/\alpha)$ identical clusters in the data. So what we will actually show is a *list-decoding* result—that it is possible to output $\mathcal{O}(1/\alpha)$ "candidates" $\hat{\mu}_l$ such that one of them is close to the mean of S:

Proposition 7.2. *Suppose that a set of points $\tilde{S} = \{x_1, \ldots, x_n\}$ contains a $(\rho, \alpha/4)$-resilient set S with mean μ. Then if $|S| \geq \alpha n$ (even if $\alpha < \frac{1}{2}$), it is possible to output $m \leq \frac{2}{\alpha}$ candidates $\hat{\mu}_1, \ldots, \hat{\mu}_m$ such that $\|\hat{\mu}_j - \mu\| \leq \frac{8\rho}{\alpha}$ for some j.*

Proof. The basic intuition is that we can cover the points in $\tilde{S}$ by resilient sets $S'_1, \ldots, S'_{2/\alpha}$ of size $\frac{\alpha}{2}n$. Then by the pigeonhole principle, the resilient set S must have large overlap with at least one of the S', and hence have similar mean. This is captured in Figure 10 below.

[6] This is not quite optimal, in reality we would want to assign based on $\|x - \mu_j\|_2^2/\sigma^2 + \log \alpha_j$, but we consider this simpler assignment for simplicity.

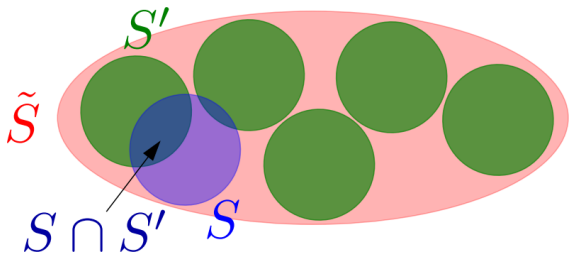

Figure 10: If we cover $\tilde{S}$ by resilient sets, at least one of the sets S' has large intersection with S.

The main difference is that S and S' may have relatively small overlap (in a roughly α-fraction of elements). We thus need to care about resilience when the subset T is small compared to S. The following lemma relates resilience on large sets to resilience on small sets:

Lemma 7.3. *For any $0 < \epsilon < 1$, a distribution/set is (ρ, ϵ)-resilient if and only if it is $(\frac{1-\epsilon}{\epsilon}\rho, 1-\epsilon)$-resilient.*

This was already proved in Appendix C as part of Lemma 2.14. Given Lemma 7.3, we can prove Proposition 7.2 with a similar triangle inequality argument to how we showed that resilient sets have small modulus of continuity. However, we now need to consider multiple resilient sets S_i rather than a single S'.

Suppose S is $(\rho, \frac{\alpha}{4})$-resilient around μ–and thus also $(\frac{4}{\alpha}\rho, 1-\frac{\alpha}{4})$-resilient by Lemma 7.3–and let $S_1, \dots, S_m$ be a maximal collection of subsets of $[n]$ such that:

1. $|S_j| \geq \frac{\alpha}{2}n$ for all j.
2. S_j is $(\frac{4}{\alpha}\rho, 1-\frac{\alpha}{2})$-resilient (with mean μ_j).
3. $S_j \cap S_{j'} = \emptyset$ for all $j \neq j'$.

Clearly $m \leq \frac{2}{\alpha}$. We claim that S has large intersection with at least one of the S_j and hence μ_j is close to μ. By maximality of the collection $\{S_j\}_{j=1}^m$, it must be that $S_0 = S \backslash (S_1 \cup \dots \cup S_m)$ cannot be added to the collection. First suppose that $|S_0| \geq \frac{\alpha}{2}n$. Then S_0 is $(\frac{4}{\alpha}\rho, 1-\frac{\alpha}{2})$-resilient (because any subset of $\frac{\alpha}{2}|S_0|$ points in S_0 is a subset of at least $\frac{\alpha}{4}|S|$ points in S). This contradicts the maximality of $\{S_j\}_{j=1}^m$, so we must have $|S_0| < \frac{\alpha}{2}n$.

Now, this implies that $|S \cap (S_1 \cup \dots \cup S_m)| \geq \frac{\alpha}{2}n$, so by pigeonhole we must have $|S \cap S_j| \geq \frac{\alpha}{2}|S_j|$ for some j. Letting $T = S \cap S_j$ as before, we find that $|T| \geq \frac{\alpha}{2}|S_j| \geq \frac{\alpha}{4}|S|$ and hence by resilience of S_j and S we have $\|\mu - \mu_j\| \leq 2 \cdot (\frac{4}{\alpha}\rho) = \frac{8}{\alpha}\rho$ by the same triangle inequality argument as before. □

Better bounds for well-separated clusters. Proposition 7.2 is powerful because it holds under very minimal conditions (we do not need to assume anything about separation of clusters or even about any of the clusters other than the one we are estimating). However, its guarantees are also minimal—we only know that we get approximate parameter recovery in the list-decoding model, and cannot say anything about cluster recovery. We next obtain a stronger bound assuming that the data can actually be separated into clusters (with a small fraction of outliers) and that the means are well-separated. This stronger result both gives cluster recovery, and gives better bounds for parameter recovery:

Proposition 7.4. *Suppose that a set of points $\{x_1, \dots, x_n\}$ can be partitioned into k sets $C_1, \dots, C_k$ of size $\alpha_1 n, \dots, \alpha_k n$, together with a fraction ϵn of outliers ($\epsilon = 1 - (\alpha_1 + \dots + \alpha_k)$), where $2\epsilon \leq \alpha = \min_{k=1}^k \alpha_j$. Further suppose that*

- *Each cluster is (ρ_1, ϵ)-resilient and $(\rho_2, 2\epsilon/\alpha)$-resilient.*
- *The means are well-separated: $\Delta > \frac{4\rho}{\epsilon}$ where $\Delta = \min_{j \neq j'} \|\mu_j - \mu_{j'}\|_2$.*

Then we can output clusters $\hat{C}_1, \dots, \hat{C}_k$ such that:

- $|C_j \triangle \hat{C}_j| \leq \mathcal{O}(\epsilon/\alpha)|C_j|$ *(cluster recovery)*

- *The mean $\hat{\mu}_j$ of $\hat{C}_j$ satisfies $\|\hat{\mu}_j - \mu_j\|_2 \leq 2\rho_2$ (parameter recovery).*

Proof. We will construct a covering by resilient sets as before, but this time make use of the fact that we know the data can be approximately partitioned into clusters. Specifically, let $S_1, \ldots, S_k$ be a collection of k sets such that:

- $|S_l| \geq \alpha n$
- The S_l are disjoint and contain all but ϵn points.
- Each S_l is (ρ_1, ϵ)-resilient.

We know that such a collection exists because we can take the C_j themselves. Now call a set S "j-like" if it contains at least $\alpha_j(\epsilon/\alpha)|S|$ points from C_j. We claim that each S_l is j-like for exactly one j. Indeed, by pigeonhole it must be j-like for at least one j since $\epsilon/\alpha \leq 1/2 < 1$.

In the other direction, note that if S if j-like then $S \cap C_j$ contains at least $(\alpha_j/\alpha)\epsilon$ of the points in S, and at least $(|S|/n)(\epsilon/\alpha) \geq \epsilon$ of the points in C_j. Thus by resilience of both sets, the means of both S and C_j are within $\frac{\rho_1}{\epsilon}$ of the mean of $S \cap C_j$ and hence within $\frac{2\rho_1}{\epsilon}$ of each other. In summary, $\|\mu_j - \mu_S\|_2 \leq \frac{2\rho_1}{\epsilon}$. Now if S were j-like and also j'-like, then we would have $\|\mu_j - \mu_{j'}\|_2 \leq \frac{4\rho_1}{\epsilon}$, which contradicts the separation assumption.

Since S_l is j-like for a unique j, it contains at most $(\epsilon/\alpha)|S_l|$ points from any of the other $C_{j'}$, together with at most ϵn outliers. Moreover, since the other $S_{l'}$ are not j-like, S_l is missing at most $\alpha_j(\epsilon/\alpha)n$ points from C_j. Thus $S_l \cap C_j$ is missing at most $2\epsilon/\alpha|S_l|$ points from S_l and at most $\epsilon/\alpha|C_j|$ points from C_j. By resilience their means are thus within $2\rho_2$ of each other, as claimed. □

[Lecture 25]

7.3 Efficient Clustering Under Bounded Covariance

We saw that resilience is information-theoretically sufficient for agnostic clustering, but we would also like to develop efficient algorithms for clustering. This is based on work in Kumar and Kannan (2010) and Awasthi and Sheffet (2012), although we will get a slightly slicker argument by using the machinery on resilience that we've developed so far.

As before, we will need a strong assumption than resilience. Specifically, we will assume that each cluster had bounded covariance and that the clusters are well-separated:

Theorem 7.5. *Suppose that the data points $x_1, \ldots, x_n$ can be split into k clusters $C_1, \ldots, C_k$ with sizes $\alpha_n, \cdots, \alpha_k n$ and means $\mu_1, \ldots, \mu_k$, and moreover that the following covariance and separation conditions hold:*

- $\frac{1}{|C_j|}\sum_{i \in C_j}(x_i - \mu_j)(x_i - \mu_j)^\top \preceq \sigma^2 I$ *for each cluster C_j,*
- $\Delta \geq 36\sigma/\sqrt{\alpha}$, *where* $\Delta = \min_{j \neq j'} \|\mu_j - \mu_{j'}\|_2$.

Then there is a polynomial-time algorithm outputting candidate clusters $\hat{C}_1, \ldots, \hat{C}_k$ and means $\hat{\mu}_1, \ldots, \hat{\mu}_k$ such that:

- $|C_j \triangle \hat{C}_j| = \mathcal{O}(\sigma^2/\alpha\Delta^2)$ *(cluster recovery), and*
- $\|\mu_j - \hat{\mu}_j\|_2 = \mathcal{O}(\sigma^2/\alpha\Delta)$ *(parameter recovery).*

The basic idea behind the algorithm is to project each of the points x_i onto the span of the top k singular vectors of the data matrix $X = [x_1 \ \cdots \ x_n]$. Let P_k be the projection operator onto this space. Then since the points Px_i lie in only a k-dimensional space instead of a d-dimensional space, they are substantially easier to cluster. The algorithm itself has three core steps and an optional step:

1. Project points x_i to Px_i.
2. Form initial clusters based on the Px_i.

3. Compute the means of each of these clusters.
4. Optionally run any number of steps of k-means in the original space of x_i, initialized with the computed means from the previous step.

We will provide more formal psuedocode later [NOTE: TBD]. For now, we focus on the analysis, which has two stages: (1) showing that the initial clustering from the first two steps is "nice enough", and (2) showing that this niceness is preserved by the k-means iterations in the second two steps.

Analyzing the projection. We start by analyzing the geometry of the points $P_k x_i$. The following lemma shows that the projected clusters are still well-separated and have small covariance:

Lemma 7.6. *The projected points $P_k x_i$ satisfy the covariance and separation conditions with parameters σ and $\sqrt{\Delta^2 - 4\sigma^2/\alpha} \geq 35\sigma/\sqrt{\alpha}$:*

$$\frac{1}{|C_j|}\sum_{i \in C_j}(Px_i - P\mu_j)(Px_i - P\mu_j)^\top \preceq \sigma^2 I \text{ and } \|P\mu_j - P\mu_{j'}\|_2 \geq \sqrt{\Delta^2 - 4\sigma^2/\alpha}. \tag{352}$$

In other words, the covariance condition is preserved, and separation is only decreased slightly.

Proof. The covariance condition is preserved because the covariance matrix of the projected points for cluster j is $P_k \Sigma_j P_k$, where Σ_j is the un-projected covariance matrix. This evidently has smaller singular values than Σ_k.

The separation condition requires more detailed analysis. We start by showing that there is not much in the orthogonal component $(I - P_k)x_i$. Indeed, we have that the top singular value of $(I - P_k)x_i$ is at most σ:

$$S = \frac{1}{n}\sum_{i=1}^{n}((I - P_k)x_i)((I - P_k)x_i)^\top \preceq \sigma^2 I \tag{353}$$

This is because P_k minimizes this top singular value among all k-dimensional projection matrices, and if we take the projection Q_k onto the space spanned by the means $\mu_1, \ldots, \mu_k$, we have

$$\frac{1}{n}\sum_{i=1}^{n}((I - Q_k)x_i)((I - Q_k)x_j)^\top = \sum_{j=1}^{k}\frac{\alpha_j}{|C_j|}\sum_{i \in C_j}((I - Q_k)x_i)((I - Q_k)x_i)^\top \tag{354}$$

$$= \sum_{j=1}^{k}\frac{\alpha_j}{|C_j|}\sum_{i \in C_j}((I - Q_k)(x_i - \mu_j))((I - Q_k)(x_i - \mu_j))^\top \tag{355}$$

$$\preceq \sum_{j=1}^{k}\frac{\alpha_j}{|C_j|}\sum_{i \in C_j}(x_i - \mu_j)(x_i - \mu_j)^\top \preceq \sum_{j=1}^{k}\alpha_j\sigma^2 I = \sigma^2 I. \tag{356}$$

Given this, we know that the projections $(I - P_k)\mu_j$ must be small, since otherwise we have

$$v^\top S v = \frac{1}{n}\sum_{i=1}^{n}\langle (I - P_k)x_i, v\rangle^2 \tag{357}$$

$$\geq \frac{\alpha_j}{|C_j|}\sum_{i \in C_j}\langle (I - P_k)x_i, v\rangle^2 \tag{358}$$

$$\geq \alpha_j\Big\langle \frac{1}{|C_j|}\sum_{i \in C_j}(I - P_k)x_i, v\Big\rangle^2 \tag{359}$$

$$= \alpha_j\langle (I - P_k)\mu_j, v\rangle^2. \tag{360}$$

Consequently $\langle (I - P_k)\mu_j, v\rangle^2 \leq \sigma^2/\alpha_j$ and hence (taking v to align with $(I - P_k)\mu_j$) we have $\|(I - P_k)\mu_j\|_2 \leq \sigma/\sqrt{\alpha_j}$. In particular $\|(I - P_k)(\mu_j - \mu_{j'})\|_2 \leq 2\sigma/\sqrt{\alpha}$.

Now, by the Pythagorean theorem we have

$$\|P_k(\mu_j - \mu_{j'})\|_2^2 = \|\mu_j - \mu_{j'}\|_2^2 - \|(I - P_k)(\mu_j - \mu_{j'})\|_2^2 \geq \Delta^2 - 4\sigma^2/\alpha, \tag{361}$$

and hence the projected means are separated by at least $\sqrt{\Delta^2 - 4\sigma^2/\alpha}$, as was to be shown. □

Analyzing the initial clustering. We now analyze the initial clustering. Call a point i a *proto-center* if there are at least $\frac{\alpha}{2}n$ projected points within distance $3\sigma\sqrt{k}$ of $P_k x_i$, and call the set of these nearby points the associated *proto-cluster*.

We will show that the proto-clusters are nearly pure (have few points not from C_j) using a similar argument as when we analyzed resilient clustering. As before, call a proto-cluster *j-like* if there are at least $\frac{\alpha_j \alpha}{4}n$ points from C_j in the proto-cluster.

Lemma 7.7. *Each proto-cluster is j-like for exactly one j.*

Proof. We know that it is j-like for at least one j by the Pigeonhole principle (if not, then the proto-cluster has at most $\frac{\alpha}{4}n$ points in total, contradicting its size of at least $\frac{\alpha}{2}n$). So suppose for the sake of contradiction that it is both j-like and j'-like. By resilience, the mean of the points from C_j is at most $2\sigma/\sqrt{\alpha}$ away from $P_k\mu_j$, and similarly the mean of the points from $C_{j'}$ is at most $2\sigma/\sqrt{\alpha}$ away from $P_k\mu_{j'}$. Since the cluster has radius $3\sigma\sqrt{k} \le 3\sigma/\sqrt{\alpha}$, this implies that $\|P_k(\mu_j - \mu_{j'})\|_2 \le 10\sigma/\sqrt{\alpha}$, contradicting the separation condition for the projected means. Thus no proto-cluster can be j-like for multiple j, which proves the lemma. □

Now since each proto-cluster is j-like for exactly one j, at least half of the points must come from that proto-cluster.

At this point we are essentially done if all we care about is constructing an efficient algorithm for cluster recovery (but not parameter recovery), since if we just extend each proto-cluster by $\mathcal{O}(\sigma)$ we are guaranteed to contain almost all of the points in its corresponding cluster, while still containing very few points from any other cluster (assuming the data are well-separated). However, parameter recovery is a bit trickier because we need to make sure that the small number of points from other clusters don't mess up the mean of the cluster. The difficulty is that while we have control over the projected distances, and can recover the projected centers $P_k\mu_j$ well, we need to somehow get back to the original centers μ_j.

The key here is that for each proto-cluster, the Px_i are all close to each other, and the missing component $(I - P_k)x_i$ has bounded covariance. Together, these imply that the proto-cluster is *resilient*—deleting an ϵ-fraction of points can change the mean by at most $\mathcal{O}(\sigma\epsilon)$ in the P_k component, and $\mathcal{O}(\sigma\sqrt{\epsilon})$ in the $(I - P_k)$ component. In fact, we have:

Lemma 7.8. *Let B be a proto-cluster with mean ν. Then*

$$\frac{1}{|B|}\sum_{i\in B}(x_i - \nu)(x_i - \nu)^\top \preceq 11\sigma^2/\alpha. \tag{362}$$

In particular, if B is j-like then we have $\|\mu_j - \nu\|_2 \le 9\sigma/\sqrt{\alpha}$.

Proof. The covariance bound is because the covariance of the x_i are bounded in norm by at most $3\sigma\sqrt{k}$ in the P_k component and hence can contribute at most $9\sigma^2 k \le 9\sigma^2/\alpha$ to the covariance, while we get an additional $2\sigma^2/\alpha$ in an orthogonal direction because the overall second moment of the $(I - P_k)x_i$ is σ^2 and the $i \in B$ contribute to at least an $\frac{\alpha}{2}$ fraction of that.

Now, this implies that B is resilient, while we already have that C_j is resilient. Since $B \cap C_j$ contains at least half the points in both B and C_j, this gives that their means are close–within distance $2(\sqrt{11}+1)\sigma/\sqrt{\alpha} < 9\sigma/\sqrt{\alpha}$. □

Analyzing k-means. We next show that k-means iterations preserve certain important invariants. We will call the assigned means $\hat{\mu}_j$ *R-close* if $\|\hat{\mu}_j - \mu_j\|_2 \le R$ for all j, and we will call the assigned clusters $\hat{C}_j$ *ϵ-close* if $|C_j \triangle \hat{C}_j| \le \epsilon|C_j|$ for all j. We will show that if the means are R-close then the clusters are ϵ-close for some $\epsilon = f(R)$, and that the resulting new means are then $g(R)$-close. If R is small enough then we will also have $g(R) < R$ so that we obtain an invariant.

Let $\Delta_{jj'} = \|\mu_j - \mu_{j'}\|_2$, so $\Delta_{jj'} \ge \Delta$. We will show that if the $\hat{\mu}_j$ are R-close, then few points in C_j can end up in $\hat{C}_{j'}$. Indeed, if x_i ends up in $\hat{C}_{j'}$ then we must have $\|x_i - \hat{\mu}_{j'}\|_2^2 \le \|x_i - \hat{\mu}_j\|_2^2$, which after some re-arrangement yields

$$\langle x_i - \hat{\mu}_j, \hat{\mu}_{j'} - \hat{\mu}_j\rangle \ge \frac{1}{4}\langle \hat{\mu}_{j'} - \hat{\mu}_j, \hat{\mu}_{j'} - \hat{\mu}_j\rangle. \tag{363}$$

Applying the covariance bound and Chebyshev's inequality along the vector $v = \hat{\mu}_{j'} - \hat{\mu}_j$, we see that the fraction of points in C_j that end up in $\hat{C}_{j'}$ is at most $\frac{4\sigma^2}{\|\hat{\mu}_j - \hat{\mu}_{j'}\|_2^2} \le \frac{4\sigma^2}{(\Delta_{jj'} - 2R)^2} \le \frac{4\sigma^2}{(\Delta - 2R)^2}$. In total this means that at most $\frac{4\sigma^2 n}{(\Delta-2R)^2}$ points from other clusters end up in $\hat{C}_j$, while at most $\frac{4k\sigma^2|C_j|}{(\Delta-2R)^2}$ points from C_j end up in other clusters. Thus we have $\epsilon \le \frac{4k\sigma^2}{(\Delta-2R)^2} + \frac{4\sigma^2}{\alpha(\Delta-2R)^2} \le \frac{8\sigma^2}{\alpha(\Delta-2R)^2}$, so we can take

$$f(R) = \frac{8\sigma^2}{\alpha(\Delta - 2R)^2}. \tag{364}$$

Now suppose that $\gamma_{jj'}|C_j|$ points in C_j are assigned to $\hat{C}_{j'}$, where we must have $\gamma_{jj'} \le \frac{4\sigma^2}{(\Delta_{jj'}-2R)^2}$. By resilience, the mean of these points is within $\sigma/\sqrt{\gamma_{jj'}}$ of μ_j and hence within $\Delta_{jj'} + \sigma/\sqrt{\gamma_{jj'}}$ of $\mu_{j'}$. In total, then, these points can shift the mean $\hat{\mu}_{j'}$ by at most

$$\frac{\gamma_{jj'}\alpha_j n(\Delta_{jj'} + \sigma/\sqrt{\gamma_{jj'}})}{\frac{1}{2}\alpha n} \le \frac{2\alpha_j}{\alpha}\Big(\frac{4\sigma^2\Delta_{jj'}}{(\Delta_{jj'} - 2R)^2} + \frac{2\sigma^2}{\Delta_{jj'} - 2R}\Big) \le \frac{4\alpha_j}{\alpha}\Big(\frac{2\sigma^2\Delta}{(\Delta - 2R)^2} + \frac{\sigma^2}{\Delta - 2R}\Big). \tag{365}$$

At the same time, the $\frac{4k\sigma^2}{(\Delta-2R)^2}$ fraction of points that are missing from $C_{j'}$ can change its mean by at most $\frac{4\sigma^2\sqrt{k}}{\Delta-2R}$. Thus in total we have

$$\|\hat{\mu}_{j'} - \mu_{j'}\|_2 \le \frac{4\sigma^2}{\Delta - 2R} \cdot \Big(\sqrt{k} + \frac{1}{\alpha} + \frac{2\Delta}{\alpha(\Delta - 2R)}\Big) \le \frac{8\sigma^2(\Delta - R)}{\alpha(\Delta - 2R)^2}. \tag{366}$$

In particular we can take $g(R) = \frac{8\sigma^2(\Delta-R)}{\alpha(\Delta-2R)^2}$.

As long as $R \le \Delta/4$ we have $g(R) \le \frac{24\sigma^2}{\alpha\Delta}$ and $f(R) \le \frac{32\sigma^2}{\alpha\Delta^2}$, as claimed. Since our initial R is $9\sigma/\sqrt{\alpha}$, this works as long as $\Delta \ge 36\sigma/\sqrt{\alpha}$, which completes the proof.

References

Radosław Adamczak and Paweł Wolff. Concentration inequalities for non-lipschitz functions with bounded derivatives of higher order. *Probability Theory and Related Fields*, 162(3-4):531–586, 2015.

Noga Alon and Assaf Naor. Approximating the cut-norm via grothendieck's inequality. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 72–80. ACM, 2004.

Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *arXiv preprint arXiv:1802.00420*, 2018.

Pranjal Awasthi and Or Sheffet. Improved spectral-norm bounds for clustering. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 37–49. Springer, 2012.

Dominique Bakry and Michel Émery. Diffusions hypercontractives. In *Séminaire de Probabilités XIX 1983/84*, pages 177–206. Springer, 1985.

Jean-Baptiste Bardet, Nathaël Gozlan, Florent Malrieu, Pierre-André Zitt, et al. Functional inequalities for gaussian convolutions of compactly supported measures: explicit bounds and dimension dependence. *Bernoulli*, 24(1): 333–353, 2018.

Stéphane Boucheron, Gábor Lugosi, and Olivier Bousquet. Concentration inequalities. In *Summer School on Machine Learning*, pages 208–240. Springer, 2003.

Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.

Nicholas Carlini and David Wagner. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pages 3–14. ACM, 2017.

Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, and Aleksander Madry. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019.

David L Donoho and Richard C Liu. The" automatic" robustness of minimum distance functionals. *The Annals of Statistics*, 16(2):552–586, 1988.

Krishnamurthy Dvijotham, Sven Gowal, Robert Stanforth, Relja Arandjelovic, Brendan O'Donoghue, Jonathan Uesato, and Pushmeet Kohli. Training verified learners with learned verifiers. *arXiv preprint arXiv:1805.10265*, 2018.

Daniel Kang, Yi Sun, Dan Hendrycks, Tom Brown, and Jacob Steinhardt. Testing robustness against unforeseen adversaries. *arXiv preprint arXiv:1908.08016*, 2019.

Amit Kumar and Ravindran Kannan. Clustering with spectral norm and the k-means algorithm. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 299–308. IEEE, 2010.

Rafał Latała. Estimates of moments and tails of gaussian chaoses. *The Annals of Probability*, 34(6):2315–2331, 2006.

Michel Ledoux. *The concentration of measure phenomenon*. Number 89. American Mathematical Soc., 2001.

Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *arXiv preprint arXiv:1712.04912*, 2017.

Aditi Raghunathan, Jacob Steinhardt, and Percy S Liang. Semidefinite relaxations for certifying robustness to adversarial examples. In *Advances in Neural Information Processing Systems*, pages 10877–10887, 2018.

Omar Rivasplata. Subgaussian random variables: An expository note. 2012.

Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.

Eric Wong and J Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. *arXiv preprint arXiv:1711.00851*, 2017.

A Properties of Statistical Discrepancies

A.1 Total variation distance

A.2 Wasserstein distance

B Concentration Inequalities

B.1 Proof of Chebyshev's inequality (Lemma 2.1)

Let $\mathbb{I}[E]$ denote the indicator that E occurs. Then we have

$$|\mathbb{E}_{X\sim p}[X \mid E] - \mu| = |\mathbb{E}_{X\sim p}[(X-\mu)\mathbb{I}[E]]|/\mathbb{P}[E] \tag{367}$$

$$\leq \sqrt{\mathbb{E}_{X\sim p}[(X-\mu)^2] \cdot \mathbb{E}_{X\sim p}[\mathbb{I}[E]^2]}/\mathbb{P}[E] \tag{368}$$

$$\leq \sqrt{\sigma^2 \cdot \mathbb{P}[E]}/\mathbb{P}[E] = \sigma/\sqrt{\mathbb{P}[E]}. \tag{369}$$

In particular, if we let E_0 be the event that $X \geq \mu + \sigma/\sqrt{\delta}$, we get that $\sigma/\sqrt{\delta} \leq \sigma/\sqrt{\mathbb{P}[E_0]}$, and hence $\mathbb{P}[E_0] \leq \delta$, which proves the first part of the lemma.

For the second part, if $\mathbb{P}[E] \leq \frac{1}{2}$ then (369) already implies the desired result since $\sigma/\sqrt{\delta} \leq \sigma\sqrt{2(1-\delta)/\delta}$ when $\delta \leq \frac{1}{2}$. If $\mathbb{P}[E] \geq \frac{1}{2}$, then consider the same argument applied to $\neg E$ (the event that E does not occur). We get

$$|\mathbb{E}_{X\sim p}[X \mid E] - \mu| = \frac{1-\mathbb{P}[E]}{\mathbb{P}[E]}|\mathbb{E}_{X\sim p}[X \mid \neg E] - \mu| \tag{370}$$

$$\leq \frac{1-\mathbb{P}[E]}{\mathbb{P}[E]} \cdot \sigma/\sqrt{1-\mathbb{P}[E]}. \tag{371}$$

Again the result follows since $\sigma\sqrt{1-\delta}/\delta \leq \sigma\sqrt{2(1-\delta)/\delta}$ when $\delta \geq \frac{1}{2}$.

B.2 Proof of d-dimensional Chebyshev's inequality (Lemma 2.8)

C Proof of Lemma 2.14

Since (ρ, ϵ)-resilience is equivalent to $(\frac{1-\epsilon}{\epsilon}\rho, 1-\epsilon)$-resilience, it suffices to show that $(1-\epsilon, \frac{1-\epsilon}{\epsilon}\rho)$-resilience is equivalent to (18). Suppose that E is an event with probability ϵ, and let v be such that $\|v\|_* = 1$ and $\langle \mathbb{E}[X-\mu \mid E], v\rangle = \|\mathbb{E}[X-\mu \mid E]\|$. Then we have

$$\begin{aligned}
\|\mathbb{E}[X-\mu \mid E]\| &= \langle \mathbb{E}[X-\mu \mid E], v\rangle && (372)\\
&= \langle \mathbb{E}[\langle X-\mu, v\rangle \mid E] && (373)\\
&\overset{(i)}{\leq} \mathbb{E}[\langle X-\mu, v\rangle \mid \langle X-\mu, v\rangle \geq \tau_\epsilon(v)] && (374)\\
&\overset{(18)}{\leq} \frac{1-\epsilon}{\epsilon}\rho. && (375)
\end{aligned}$$

Here (i) is because $\langle X-\mu, v\rangle$ is at least as large for the ϵ-quantile as for any other event E of probability ϵ. This shows that (18) implies $(1-\epsilon, \frac{1-\epsilon}{\epsilon}\rho)$-resilience. For the other direction, given any v let E_v denote the event that $\langle X-\mu, v\rangle \geq \tau_\epsilon(v)$. Then E_v has probability ϵ and hence

$$\begin{aligned}
\mathbb{E}[\langle X-\mu, v\rangle \mid \langle X-\mu, v\rangle \geq \tau_\epsilon(v)] &= \mathbb{E}[\langle X-\mu, v\rangle \mid E_v] && (376)\\
&= \langle \mathbb{E}[X-\mu \mid E_v], v\rangle && (377)\\
&\overset{(ii)}{\leq} \|\mathbb{E}[X-\mu \mid E_v]\| && (378)\\
&\overset{(iii)}{\leq} \frac{1-\epsilon}{\epsilon}\rho, && (379)
\end{aligned}$$

where (ii) is Hölder's inequality and (iii) invokes resilience. Therefore, resilience implies (18), so the two properties are equivalent, as claimed.

D Proof of Lemma 2.15

Let E_+ be the event that $\langle x_i - \mu, v\rangle$ is positive, and E_- the event that it is non-negative. Then $\mathbb{P}[E_+]+\mathbb{P}[E_-] = 1$, so at least one of E_+ and E_- has probablity at least $\frac{1}{2}$. Without loss of generality assume it is E_+. Then we have

$$\begin{aligned}
\mathbb{E}_{x\sim p}[|\langle x-\mu, v\rangle|] &= 2\mathbb{E}_{x\sim p}[\max(\langle x-\mu, v\rangle, 0)] && (380)\\
&= 2\mathbb{P}[E_+]\mathbb{E}_{x\sim p}[\langle x-\mu, v\rangle \mid E_+] && (381)\\
&\leq 2\mathbb{P}[E_+]\|\mathbb{E}_{x\sim p}[x-\mu \mid E_+]\| \leq 2\rho, && (382)
\end{aligned}$$

where the last step invokes resilience applies to E_+ together with $\mathbb{P}[E_+] \leq 1$. Conversely, if p has bounded 1st moments then

$$\begin{aligned}
\mathbb{E}[\langle X-\mu, v\rangle \mid \langle X-\mu, v\rangle \geq \tau_{1/2}(v)] &\leq \mathbb{E}[|\langle X-\mu, v\rangle|]/\mathbb{P}[\langle X-\mu, v\rangle \geq \tau_{1/2}(v)] && (383)\\
&= 2\mathbb{E}[|\langle X-\mu, v\rangle|] \leq 2\rho, && (384)
\end{aligned}$$

so p is $(2\rho, \frac{1}{2})$-resilient by Lemma 2.14.