

STATISTICAL ESTIMATION: FROM DENOISING TO SPARSE REGRESSION AND HIDDEN CLIQUES

ERIC W. TRAMEL, SANTHOSH KUMAR, ANDREI GIURGIU, AND ANDREA MONTANARI

ABSTRACT. These notes review six lectures given by Prof. Andrea Montanari on the topic of statistical estimation for linear models. The first two lectures cover the principles of signal recovery from linear measurements in terms of minimax risk. Subsequent lectures demonstrate the application of these principles to several practical problems in science and engineering. Specifically, these topics include denoising of error-laden signals, recovery of compressively sensed signals, reconstruction of low-rank matrices, and also the discovery of hidden cliques within large networks.

These are notes from the lecture of Andrea Montanari given at the autumn school “Statistical Physics, Optimization, Inference, and Message-Passing Algorithms”, that took place in Les Houches, France from Monday September 30th, 2013, till Friday October 11th, 2013. The school was organized by Florent Krzakala from UPMC & ENS Paris, Federico Ricci-Tersenghi from La Sapienza Roma, Lenka Zdeborová from CEA Saclay & CNRS, and Riccardo Zecchina from Politecnico Torino.

A.M. was partially supported by the NSF grant CCF-1319979 and the grants AFOSR/DARPA FA9550-12-1-0411 and FA9550-13-1-0036.

E.W.T. was supported by the ERC under the European Union’s 7th Framework Programme Grant Agreement 307087-SPARCS.

CONTENTS

Preface	2
1. Statistical estimation and linear models	3
1.1. Statistical estimation	3
1.2. Denoising	6
1.3. Least Squares (LS) estimation	6
1.4. Evaluating the estimator	10
2. Nonlinear denoising and sparsity	14
2.1. Minimax risk	14
2.2. Approximation error and the bias-variance tradeoff	15
2.3. Wavelet expansions	19
3. Denoising with thresholding	22
3.1. An equivalent analysis: Estimating a random scalar	26
4. Sparse regression	28
4.1. Motivation	28
4.2. The LASSO	29
4.3. Behavior of the LASSO under Restricted Isometry Property	33
5. Random designs and Approximate Message Passing	35
5.1. Message Passing algorithms	36
5.2. Analysis of AMP and the LASSO	38
6. The hidden clique problem	42
6.1. An iterative thresholding approach	47
6.2. A message passing algorithm	49
6.3. Analysis and optimal choice of $f_i(\cdot)$	50
References	52

PREFACE

These lectures provide a gentle introduction to some modern topics in high-dimensional statistics, statistical learning and signal processing, for an audience without any previous background in these areas. The point of view we take is to connect the recent advances to basic background in statistics (estimation, regression and the bias-variance trade-off), and to classical –although non-elementary– developments (sparse estimation and wavelet denoising).

The first three sections will cover these basic and classical topics. We will then cover more recent research, and discuss sparse linear regression in Section 4.3.1, and its analysis for random designs in Section 5.2. Finally, in Section 6.3 we discuss an intriguing example of a class of problems whereby sparse and low-rank structures have to be exploited simultaneously.

Needless to say, the selection of topics presented here is very partial. The reader interested in a deeper understanding can choose from a number of excellent options for further study. Very readable introductions to the fundamentals of statistical estimation can be found in the books by Wasserman [1, 2]. More advanced references (with a focus on high-dimensional and non-parametric settings) are the monographs by Johnstone [3] and Tsybakov [4]. The recent book by Bühlmann and van de Geer

[5] provides a useful survey of recent research in high-dimensional statistics. For the last part of these notes, dedicated to most recent research topics, we will provide references to specific papers.

1. STATISTICAL ESTIMATION AND LINEAR MODELS

1.1. Statistical estimation. The general problem of statistical estimation is the one of estimating an unknown object from noisy observations. To be concrete, we can consider the model

$$y = f(\theta; \text{noise}), \quad (1)$$

where y is a set of observations, θ is the unknown object, for instance a vector, a set of parameters, or a function. Finally, $f(\cdot; \text{noise})$ is an observation model which links together the observations and the unknown parameters which we wish to estimate. Observations are corrupted by random noise according to this model. The objective is to produce an estimation $\hat{\theta} = \hat{\theta}(y)$ that is accurate under some metric. The estimation of θ from y is commonly aided by some hypothesis about the structure, or behavior, of θ . Several examples are described below.

Statistical estimation can be regarded as a subfield of statistics, and lies at the core of a number of areas of science and engineering, including data mining, signal processing, and inverse problems. Each of these disciplines provides some information on how to model data acquisition, computation, and how best to exploit the hidden structure of the model of interest. Numerous techniques and algorithms have been developed over a long period of time, and they often differ in the assumptions and the objectives that they try to achieve. As an example, a few major distinctions to keep in mind are the following.

Parametric versus non-parametric: In parametric estimation, stringent assumptions are made about the unknown object, hence reducing θ to be determined by a small set of parameters. In contrast, non-parametric estimation strives to make minimal modeling assumptions, resulting in θ being an high-dimensional or infinite-dimensional object (for instance, a function).

Bayesian versus frequentist: The Bayesian approach assumes θ to be a random variable as well, whose ‘prior’ distribution plays an obviously important role. From a frequentist point of view, θ is instead an arbitrary point in a set of possibilities. In these lectures we shall mainly follow the frequentist point of view, but we stress that the two are in fact closely related.

Statistical efficiency versus computational efficiency: Within classical estimation theory, a specific estimator $\hat{\theta}$ is mainly evaluated in terms of its accuracy: How close (or far) is $\hat{\theta}(y)$ to θ for typical realizations of the noise? We can broadly refer to this figure of merit as to ‘statistical efficiency.’

Within modern applications, computational efficiency has arisen as a second central concern. Indeed θ is often high-dimensional: it is not uncommon to fit models with millions of parameters. The amounts of observations has grown in parallel. It becomes therefore crucial to devise estimators whose complexity scales gently with the dimensions, and with the amount of data.

We next discuss informally a few motivating examples.

1.1.1. *Example 1: Exploration seismology.* Large scale statistical estimation plays a key role in the field of exploration seismology. This technique uses seismic measurements on the earth surface to reconstruct geological structures, composition and density field of a geological substrates in [6]. Measurements are acquired, generally, by sending some known seismic wave through the ground, perhaps through a controlled explosive detonation, and measuring the response at multiple spatially dispersed sensors.

Below is a simple dictionary that points at the various elements of the model (1) in this example.

Exploration Seismology	
y	seismographic measurements
θ	earth density field
<i>Hypothesis</i>	smooth density field

The function $f(\dots)$ in Eq. (1) expresses the outcome of the seismographic measurements, given a certain density field θ and a certain source of seismic waves (left implicit since it is known). While this relation is of course complex, and ultimately determined by the physics of wave propagation, it is in principle perfectly known.

Because of the desired resolution of the recovered earth density field, this statistical estimation problem is often ill-posed, as sampling is severely limited by the cost of generating the source signal and the distribution and set-up of the receivers. Resolution can be substantially improved by using some structural insights into the nature of the earth density field. For instance, one can exploit the fact that this is mostly smooth with the exception of some discontinuity surfaces.





FIGURE 1. A recovered earth density field, from [6].

1.1.2. *Example 2: Hidden structure in networks.* Many modern data sets are relational, i.e. they express pairwise relations within a set of objects. This is the case in social networks, communication networks, unsupervised learning and so on.

In the simplest case, for each pair of nodes in a network, we know whether they are connected or not. Finding a hidden structure in such a network is a recurring problem with these datasets. A highly idealized but nevertheless very interesting problem requires to find a highly connected subgraph in a otherwise random graph.


Hidden Network Structure	
y	large network
θ	hidden subset of nodes
<i>Hypothesis</i>	hidden network is highly connected

From Figure 2, it is apparent that the discovery of such networks can be a difficult task.



(a) Subgraph easily visible.

(b) Subgraph hidden.



(c) Subgraph revealed

FIGURE 2. The same network adjacency matrix, is shown in (a) and (b), but the nodes permuted in (b). In (c), the hidden subgraph is revealed.

1.1.3. *Example 3: Collaborative filtering.* Recommendation systems are ubiquitous in e-commerce and web services. They aim at personalizing each user’s experience through an analysis of her past behavior, and –crucially– the past behavior of similar users. The algorithmic and statistical techniques that allow to exploit this information are referred to as ‘collaborative filtering.’ Amazon, Netflix, YouTube all make intensive use of collaborative filtering technologies.

In a idealized model for collaborative filtering, each user of a e-commerce site is associated to a row of a matrix, and each product to a column. Entry $\theta_{i,j}$ in this matrix corresponds to the evaluation that user i gives of product j . A small subset of the entries is observed because of feedback provided by the users (reviews, ratings, purchasing behavior). In this setting, collaborative filtering aims at estimating the whole matrix, on the basis of noisy observations of relatively few of its entries.

While this task is generally hopeless, it is observed empirically that such data matrices are often well approximated by low-rank matrices. This corresponds to the intuition that a small number of factors (corresponding to the approximate rank) explain the opinions of many users concerning many items. The problem is then modeled as the one of estimating a low-rank matrix from noisy observations of some of its entries.

Collaborative Filtering	
y	small set of entries in a large matrix
θ	unknown entries of matrix
<i>Hypothesis</i>	matrix has a low-rank representation

A toy example of this problem is demonstrated in Figures 3-4. It can be observed that an accurate estimation of the original matrix is possible even when very few of its coefficients are known.

1.2. Denoising. We will begin by considering in greater depth a specific statistical estimation problem, known as ‘denoising.’ One the one hand, denoising is interesting, since it is a very common signal processing task: In essence, it seeks restore a signal which has been corrupted by some random process, for instance additive white-noise. On the other, it will allow us to introduce some basic concepts that will play an important role throughout these lectures. Finally, recent research by [7] has unveiled a deep and somewhat surprising connection between denoising and the rapidly developing field of compressed sensing.

To formally define the problem, we assume that the signal to be estimated is a function $t \mapsto f(t)$. Without loss of generality, we will restrict the domain of $f(t)$, $f : [0, 1] \rightarrow \mathbb{R}$. We measure n uniformly-spaced samples over the domain of f ,

$$y_i = f(i/n) + w_i, \quad (2)$$

where $i \in \{1, 2, \dots, n\}$ is the sample index, and $w_i \sim \mathcal{N}(0, \sigma^2)$ is the additive noise term. Each of y_1, \dots, y_n is a sample.

For the denoising problem, we desire to calculate the original function from the noise-corrupted observables y_i . How might we go about doing this?

1.3. Least Squares (LS) estimation. The **Least Squares Method** dates back to Gauss and Legendre [8].

A natural first idea is to parametrize the function f . For instance we can tentatively assume that it is a degree $p - 1$ polynomial

$$f(t) = \sum_{j=1}^p \theta_j t^{j-1}. \quad (3)$$

Each monomial is weighted according to coefficient θ_j for $j \in \{1, 2, \dots, p\}$, and we will collect these coefficients in a vector $\theta = (\theta_1, \theta_2, \dots, \theta_p) \in \mathbb{R}^p$. Thus, the problem of recovering $f(t)$ boils down to the recovery of the p coefficients θ_j from the set of observables, y_i . We therefore seek to find the set of coefficients which which generate a function that most closely matches the observed samples.

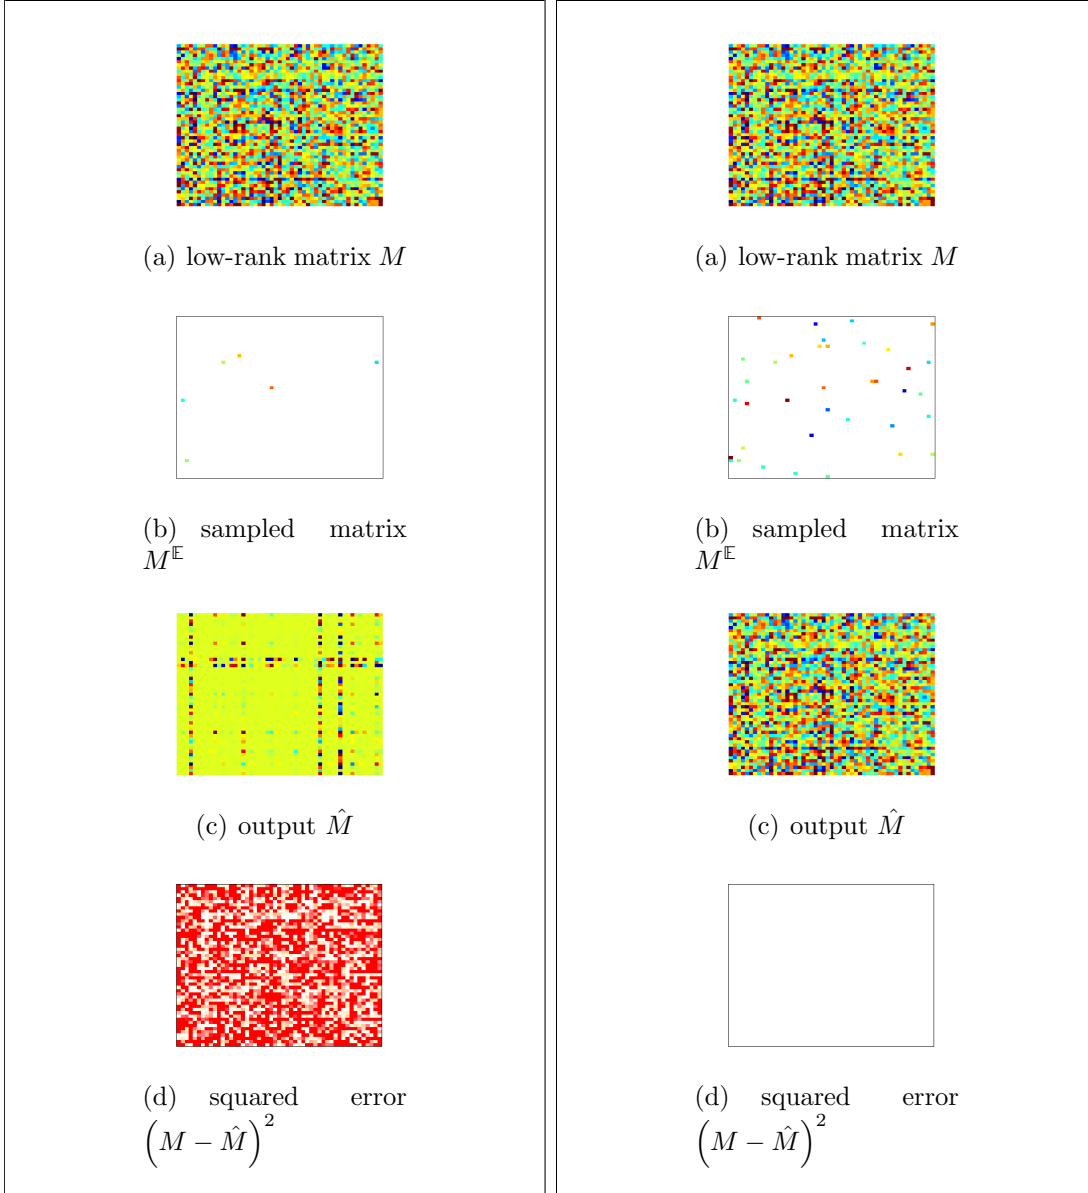


FIGURE 3. Recovering a 2000×2000 rank-8 matrix from 0.25% of its values.

FIGURE 4. Recovering a 2000×2000 rank-8 matrix from 1.75% of its values.

It is natural to set this up as an optimization problem (here RSS stands for ‘residual sum of squares’)

$$\hat{\theta}^{\text{LS}}(y) \equiv \underset{\theta}{\operatorname{argmin}} \operatorname{RSS}(\theta), \tag{4}$$

$$\operatorname{RSS}(\theta) \equiv \sum_{i=1}^n \left(y_i - \sum_{j=1}^p \theta_j \left(\frac{i}{n} \right)^{j-1} \right)^2. \tag{5}$$

Fitting a low-degree polynomial to a dataset by least squares is a very common practice, and the reader has probably tried this exercise at least once. A moment




FIGURE 5. Depiction of a discrete-time sampling of the continuous-time function, $f(x)$. Notice that the additive noise term, w_i , prevents exact knowledge of $f(i/n)$.

reflection reveals that nothing is special about the polynomials used in this procedure. In general, we can consider a set of functions $\{\varphi_1, \varphi_2, \dots, \varphi_p\}$, where

$$\varphi_j : [0, 1] \rightarrow \mathbb{R}. \quad (6)$$

Of course, the quality of our estimate depends on how well the functions $\{\varphi_j\}$ capture the behavior of the signal f . Assuming that f can be represented as a linear combination of these functions, we can rewrite our model as

$$y_i = \sum_{j=1}^p \theta_{0j} \varphi_j(i/n) + w_i. \quad (7)$$

Equivalently, if we define $\varphi : [0, 1] \rightarrow \mathbb{R}^p$ by letting $\varphi(x) = (\varphi_1(x), \varphi_2(x), \dots, \varphi_p(x))$, $\theta_0 = (\theta_{0,1}, \theta_{0,2}, \dots, \theta_{0,p})$, and denoting by $\langle a, b \rangle \equiv \sum_{i=1}^m a_i b_i$ the usual scalar product in \mathbb{R}^m , we have

$$y_i = \langle \theta_0, \varphi(i/n) \rangle + w_i. \quad (8)$$

Before continuing further, it is convenient to pass to matrix notation. Let us define a matrix $\mathbf{X} (X_{ij}) \in \mathbb{R}^{n \times p}$ whose entry i, j is given by

$$X_{ij} = \varphi_j\left(\frac{i}{n}\right), \quad (9)$$

Using this notation, and letting $y = (y_1, y_2, \dots, y_n)$, $w = (w_1, w_2, \dots, w_n)$, our model reads

$$y = X\theta_0 + w, \quad (10)$$

$w \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_n)$ (here and below \mathbf{I}_n denotes the identity matrix in n dimensions: the subscript will be dropped if clear from the context).

From (10), we see that vector of observations y is approximated as a linear combination of the columns of \mathbf{X} , each columns corresponding to one of the functions $\varphi_1, \varphi_2, \dots, \varphi_p$, evaluated on the sampling points.

This is a prototype of a very general idea in statistical learning, data mining and signal processing. Each data point x (or each point in a complicated space, e.g. a space of images) is represented by a vector in \mathbb{R}^p . This vector is constructed by evaluating p functions at x hence yielding the vector $(\varphi_1(x), \varphi_2(x), \dots, \varphi_p(x))$. Of course, the choice suitable functions $\{\varphi_j\}$ is very important and domain-specific.

The functions $\{\varphi_j\}$ (or –correspondingly– the columns of the matrix \mathbf{X}) have a variety of names. They are known as “covariates” and “predictors” in statistics, as “features” in the context of machine learning and pattern recognition. The set of features $\{\varphi_j\}$ is sometimes called a “dictionary,” and the matrix \mathbf{X} is also referred to as the “design matrix.” Finding an appropriate set of features, i.e. “featurizing”, is a problem of its own. The observed y_i are commonly referred to as the “responses” or “labels” within statistics and machine-learning, respectively. The act of finding the true set of coefficients θ_0 is known as both “regression” and “supervised learning”.

So, how do we calculate the coefficients θ_0 from y ? Going back to least squares estimation, we desire to find a set of coefficients, $\hat{\theta}$ which best match our observations. Specifically, in matrix notation (4) reads

$$\hat{\theta}^{\text{LS}} = \underset{\theta \in \mathbb{R}^p}{\operatorname{argmin}} \mathcal{L}(\theta), \quad (11)$$

where

$$\begin{aligned} \mathcal{L}(\theta) &= \frac{1}{2n} \|y - \mathbf{X}\theta\|_2^2, \\ &= \frac{1}{2n} \sum_{i=1}^n (y_i - \langle x_i, \theta \rangle)^2, \end{aligned} \quad (12)$$

with x_i the i -th row of \mathbf{X} . Here and below $\|a\|_2$ denotes the ℓ_2 -norm of vector a : $\|a\|_2^2 = \sum_i a_i^2$. The minimizer can be found by noting that

$$\nabla \mathcal{L}(\theta) = -\frac{1}{n} \mathbf{X}^\top (y - \mathbf{X}\theta), \quad (13)$$

$$\therefore \hat{\theta}^{\text{LS}} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top y. \quad (14)$$

Looking at (14), we note that an important role is played by the sample covariance matrix

$$\hat{\Sigma} = \frac{1}{n} \mathbf{X}^\top \mathbf{X}. \quad (15)$$

This is the matrix of correlations of the predictors $\{\varphi_1, \dots, \varphi_p\}$. The most immediate remark is that, for $\hat{\theta}^{\text{LS}}$ to be well defined, $\hat{\Sigma}$ needs to be invertible, which is equivalent to require $\operatorname{rank}(\mathbf{X}) = p$. This of course can only happen if the number of parameter is no larger than the number of observations: $n \leq p$. Of course, if $\hat{\Sigma}$ is invertible but is nearly-singular, then $\hat{\theta}$ will be very unstable and hence a poor estimator. A natural way to quantify the ‘goodness’ of $\hat{\Sigma}$ is through its condition number $\kappa(\hat{\Sigma})$, that is the ratio of its largest to its smallest eigenvalue: $\kappa(\hat{\Sigma}) = \lambda_{\max}(\hat{\Sigma})/\lambda_{\min}(\hat{\Sigma})$. From this point of view, an optimal design has minimal condition number $\kappa(\hat{\Sigma}) = 1$, which corresponds to \mathbf{X} to be proportional to an orthogonal matrix. In this case \mathbf{X} is called an ‘orthogonal design’ and we shall fix normalizations by assuming $\hat{\Sigma} = (\mathbf{X}^\top \mathbf{X}/n) = \mathbf{I}_p$

In functional terms, we see that the LS estimator is calculated according to the correlations between y and the predictors,

$$\widehat{\Sigma}_{jl} = \frac{1}{n} \sum_{i=1}^n \varphi_j(i/n) \varphi_l(i/n), \quad (16)$$

$$\therefore \widehat{\theta}_l^{\text{LS}} = \sum_{j=1}^p \left(\widehat{\Sigma}^{-1} \right)_{lj} \left(\frac{1}{n} \sum_{i=1}^n \varphi_j(i/n) y_i \right). \quad (17)$$

1.4. Evaluating the estimator. Now that we calculated the LS estimator for our problem, a natural question arises: is this indeed the best estimator we could use? In order to answer this question, we need a way of comparing one estimator to another.

This is normally done by considering the *risk function* associated with the estimator. If the model depends on a set of parameters $\theta \in \mathbb{R}^p$, the risk function is a function $R : \mathbb{R}^p \rightarrow \mathbb{R}$, defined by

$$\begin{aligned} R(\theta) &= \mathbb{E} \left\{ \|\widehat{\theta}(y) - \theta\|_2^2 \right\}, \\ &= \sum_{j=1}^p \mathbb{E} \left\{ (\widehat{\theta}_j(y) - \theta_j)^2 \right\}. \end{aligned} \quad (18)$$

Here expectation is taken with respect to y , distributed according to the model (10) with $\theta_0 = \theta$. Note that the ℓ_2 -distance is used to measure the estimation error.

Other measures (called ‘loss functions’) could be used as well, but we will focus on this for the sake of concreteness. We can also calculate risk over the function space and not just over the parameter space. This is also known as the ‘prediction error’:

$$\begin{aligned} R_p(\theta) &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left\{ \left(\widehat{f}(i/n) - f(i/n) \right)^2 \right\}, \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left\{ \left[\sum_{j=1}^p \mathbf{X}_{ij} (\widehat{\theta}_j - \theta_j) \right]^2 \right\}, \\ &= \frac{1}{n} \mathbb{E} \left\{ \|\mathbf{X} (\widehat{\theta} - \theta)\|_2^2 \right\}. \end{aligned} \quad (19)$$

In particular, for \mathbf{X} an orthogonal design, $R_p(\theta) = c R(\theta)$.

Let us apply this definition of risk to the LS estimator, $\widehat{\theta}^{\text{LS}}$. Returning to the signal sampling model,

$$y = \mathbf{X}\theta_0 + w. \quad (20)$$

$$\begin{aligned} \therefore \widehat{\theta}^{\text{LS}} &= (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T y, \\ &= \theta_0 + (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T w, \end{aligned} \quad (21)$$

which shows that the LS estimator will return the true parameters, θ_0 , perturbed by some amount due to noise. Now, we will calculate the risk function

$$\begin{aligned}
 R(\theta) &= \mathbb{E} \left\{ \|\widehat{\theta}^{\text{LS}}(y) - \theta_0\|_2^2 \right\}, \\
 &= \mathbb{E} \left\{ \left\| (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top w \right\|_2^2 \right\}, \\
 &= \mathbb{E} \left\{ w^\top \mathbf{X} (\mathbf{X}^\top \mathbf{X})^{-2} \mathbf{X}^\top w \right\}, \\
 &= \sigma^2 \text{Tr} \left(\mathbf{X} (\mathbf{X}^\top \mathbf{X})^{-2} \mathbf{X}^\top \right), \\
 &= \sigma^2 \text{Tr} \left((\mathbf{X}^\top \mathbf{X})^{-1} \right), \\
 &= \frac{\sigma^2 p}{n} \left[\frac{\text{Tr} \left(\widehat{\Sigma}^{-1} \right)}{p} \right], \tag{22}
 \end{aligned}$$

where we add the p term to the final result because we expect that $\frac{1}{p} \text{Tr} \left(\widehat{\Sigma}^{-1} \right)$ to be on the order one, under the assumption of near-orthonormal predictors.

To further illustrate this point, let us consider the case in which the functions $\{\varphi_j\}$ are orthonormal (more precisely, they are an orthonormal set in $L^2([0, 1])$). This means that

$$\int_0^1 \varphi_i(x) \varphi_j(x) dx = \delta_{ij}. \tag{23}$$

where δ_{ij} is 1 when $i = j$ and 0 for all $i \neq j$. For n large, this implies

$$\widehat{\Sigma}_{jl} = \frac{1}{n} \sum_{i=1}^n \varphi_j(i/n) \varphi_l(i/n) \approx \delta_{jl}, \tag{24}$$

where we assumed that the sum can be approximated by an integral. In other words, if the functions $\{\varphi_j\}$ are orthonormal, the design is nearly orthogonal, and this approximation gets better as the number of samples increases. Thus, in such good conditions, $\text{Tr} \left(\widehat{\Sigma}^{-1} \right) \approx p$. Under these conditions, we can simplify the risk function for the LS estimator in the case of orthonormal or near-orthonormal predictors to be,

$$R(\theta) \approx \frac{p\sigma^2}{n}. \tag{25}$$

This result has several interesting properties:

- The risk is proportional to the noise variance. This makes sense: the larger is noise, the worse we can estimate the function.
- It is inversely proportional to the number of samples n : the larger is the number of observations, the better we can estimate f .
- The risk is proportional to the number of parameters p . This fact can be interpreted as an over-fitting phenomenon. If we choose a large p , then our estimator will be more sensitive to noise. Conversely, we are effectively searching the right parameters in a higher-dimensional space, and a larger number of samples is required to determine it to the same accuracy.

- The risk $R(\theta)$ is independent of θ . This is closely related to the e linearity of the LS estimator.

At this point, two questions arise naturally:

- Q1. Is this the best that we can do? Is it possible to use a different estimator and decrease risk?
- Q2. What happens if the function to be estimated is function is not *exactly* given by a linear combinations of the predictors? Indeed in general, we cannot expect to have a perfect model for the signal, and any set of predictors is only approximate:

$$f(t) \neq \sum_{j=1}^p \theta_j \varphi_j(t). \quad (26)$$

To discuss these two issues, let us first change the notation for the risk function, by making explicit the dependence on the estimator $\hat{\theta}$:

$$R(\theta; \hat{\theta}) \quad \text{where} \quad \theta \in \mathbb{R}^p, \quad (27)$$

$$\hat{\theta} : \mathbb{R}^n \rightarrow \mathbb{R}^p.$$

Note that estimators, $\hat{\theta}$, are functions of $y \in \mathbb{R}^n$. For two estimators, $\hat{\theta}^1$ and $\hat{\theta}^2$, we can compare $R(\theta; \hat{\theta}^1)$ and $R(\theta; \hat{\theta}^2)$. This leads to the next, crucial question: how do we compare these two curves? For instance, in Figure 6 we sketch two cartoon risk functions $R(\theta; \hat{\theta}^1)$ and $R(\theta; \hat{\theta}^2)$. Which one is the best one? The way this question is answered has important consequences.

Note that naively, one could hope to find an estimator that is *simultaneously* the best at all points θ . Letting this ideal estimator be denoted by $\hat{\theta}^{\text{opt}}$, we would get

$$R(\theta; \hat{\theta}^{\text{opt}}) \leq R(\theta; \hat{\theta}) \quad \forall \theta, \hat{\theta}. \quad (28)$$

However, assuming the existence of such an ideal estimator leads to a contradiction. To see this, we will let the predictors be, for simplicity

$$\mathbf{X} = \sqrt{n} \mathbf{I}_n, \quad p = n, \quad (29)$$

which means our regression problem is now

$$y = \theta + \frac{w}{\sqrt{n}}, \quad w \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_n). \quad (30)$$

Note that in this case the LS estimator is simply $\hat{\theta}^{\text{LS}}(y) = y$.

Next, fix $\xi \in \mathbb{R}^p$, and consider the oblivious estimator that always returns ξ :

$$\hat{\theta}^\xi(y) = \xi. \quad (31)$$

This has the risk function

$$R(\theta; \hat{\theta}^\xi) = \|\xi - \theta\|_2^2. \quad (32)$$

If an ‘ideal’ estimator $\hat{\theta}^{\text{opt}}$ as above existed, it would beat $\hat{\theta}^\xi$, which implies in particular

$$R(\xi; \hat{\theta}^{\text{opt}}) = 0. \quad (33)$$




FIGURE 6. Comparing the risk functions of two different estimators, $\hat{\theta}^1$, $\hat{\theta}^2$ over the space of possible parameters, θ . Also shown is a risk function for some estimator which is *ideal* in the sense that it is below both both of the known estimators for all θ .




FIGURE 7. Any arbitrary trivial estimator can minimize the point-wise risk at point ξ .

Since ξ is arbitrary, this would imply that the ideal estimator has risk everywhere equal to 0, i.e. always reconstruct the true signal perfectly, independently of the noise. This is of course impossible.

One approach would be to evaluate the Bayes risk, which would compute the expected value of each risk curve dependent upon the prior distribution of the parameters, $\Pr[\theta]$. However, it is not clear in every case how one might determine this prior, and its choice can completely skew the comparison between $\hat{\theta}^1$ and $\hat{\theta}^2$.

One approach to overcome this problem is to evaluate for each risk function the corresponding ‘Bayes risk.’ This amounts to averaging $R(\theta; \hat{\theta})$ over θ , using a certain prior distribution of the parameters $P(\theta)$. Namely

$$R_B(P; \hat{\theta}) \equiv \int R(\theta; \hat{\theta}) P(d\theta). \tag{34}$$

However, it is not clear in every case how one might determine this prior. Further, the choice of $P(\theta)$ can completely skew the comparison between $\hat{\theta}^1$ and $\hat{\theta}^2$. If $P(\theta)$ is concentrated in a region in which –say– $\hat{\theta}^1$ is superior to $\hat{\theta}^2$, then $\hat{\theta}^1$ will obviously win the comparison, and viceversa.

In the next section we shall discuss the minimax approach to comparing estimators.




FIGURE 8. The risk of an unknown estimator which dominates $\hat{\theta}^{LS}$ for all θ .

2. NONLINEAR DENOISING AND SPARSITY

2.1. Minimax risk. The previous lecture discussed estimating a set of parameters, θ , given the linear model

$$y = \mathbf{X}\theta + w. \quad (35)$$

In this discussion, we stated that there exists no estimator which dominates all other possible estimators in terms of risk, $R(\theta; \hat{\theta})$. Still the question remains of how to compare two different estimators $\hat{\theta}^1, \hat{\theta}^2$.

A fruitful approach to this question is to consider the worst case risk over some region $\Omega \subseteq \mathbb{R}^p$. Formally, we define the *minimax risk* of $\hat{\theta}$ over Ω as

$$R_*(\Omega; \hat{\theta}) = \sup_{\theta \in \Omega} R(\theta; \hat{\theta}). \quad (36)$$

Such a definition of risk is useful if we have some knowledge *a priori* about the region in which the true parameters live. The minimax risk allows us to compare the maximal risk of a given estimator over the set Ω to find an estimator with minimal worst case risk. The minimax risk is also connected to the Bayes risk which we defined earlier,

$$R_B(\mathbb{P}; \hat{\theta}) \equiv \int R(\theta; \hat{\theta}) \mathbb{P}(d\theta), \quad (37)$$

$$R_*(\Omega; \hat{\theta}) = \sup_{\text{supp}(\mathbb{P}) \subseteq \Omega} R_B(\mathbb{P}; \hat{\theta}). \quad (38)$$

With this definition of minimax risk, it is easy to compute the minimax risk of least squares

$$R_*(\mathbb{R}^p; \hat{\theta}^{LS}) = \frac{p\sigma^2}{n} \left[\frac{\text{Tr}(\hat{\Sigma}^{-1})}{p} \right], \quad (39)$$

The least squares estimator is optimal in minimax sense.

Theorem 1. *The least squares estimator is minimax optimal over \mathbb{R}^p . Namely, any estimator $\hat{\theta}$ has minimax risk $R_*(\mathbb{R}^p; \hat{\theta}) \geq R_*(\mathbb{R}^p; \hat{\theta}^{LS})$.*

Proof. The proof of this result relies on the connection with Bayes risk. Consider for the sake of simplicity the case of orthogonal designs, $\hat{\Sigma} = \mathbf{I}$. It is not hard to show

that, if P is gaussian with mean 0 and covariance $c^2 \mathbf{I}_p$, then

$$\inf_{\hat{\theta}} R_B(P; \hat{\theta}) = \frac{c^2 \sigma^2}{c^2 + \sigma^2}. \quad (40)$$

Hence, for any estimator $\hat{\theta}$

$$\sup_{\theta \in \mathbb{R}^p} R(\theta; \hat{\theta}) \geq R_B(P; \hat{\theta}) \geq \frac{c^2 \sigma^2}{c^2 + \sigma^2}. \quad (41)$$

Since c is arbitrary, we can let $c \rightarrow \infty$, whence

$$R_*(\mathbb{R}^p; \hat{\theta}) \geq \sigma^2. \quad (42)$$

A full treatment of a more general result can be found, for instance in [2, Chapter 7]. \square

A last caveat. One might suspect –on the grounds of the last theorem– that least squares estimation is optimal ‘everywhere’ in \mathbb{R}^p . This was indeed common belief among statisticians until the surprising discovery of the ‘Stein phenomenon’ in the early sixties [9]. In a nutshell, for $p \geq 3$ there exist estimators that have risk $R(\theta; \hat{\theta}) < R(\theta; \hat{\theta}^{\text{LS}})$ strictly for every $\theta \in \mathbb{R}^p$! (The gap vanishes as $\theta \rightarrow \infty$.) We refer to [2, Chapter 7] for further background on this.

2.2. Approximation error and the bias-variance tradeoff. Until now we have assumed that the unknown function $f(t)$, could be exactly represented by the set of predictors, corresponding to columns of \mathbf{X} . How is our ability to estimate the parameters set θ , and thus $f(t)$, affected when this assumption is violated?

In order to study this case, we assume that we are given an infinite sequence of predictors $\{\varphi_j\}_{j \geq 1}$, and use only the first J to estimate f . For any fixed J , $f(t)$ can be approximated as a linear combination of the first J predictors, plus an error term which is dependent upon J

$$f(t) = \sum_{j=1}^J \theta_j \varphi_j(t) + \Delta_J(t). \quad (43)$$

For a complete set $\{\varphi_j\}$, we can ensure $\lim_{J \rightarrow \infty} \|\Delta_J\| = 0$ in a suitable norm. This can be formalized by assuming $\{\varphi_j\}_{j \geq 1}$ to be an orthonormal basis in the Hilbert space $L^2([0, 1])$ and the above to be the orthonormal decomposition. In particular, the remainder will be orthogonal to the expansion,

$$\int_0^1 \Delta_J(t) \varphi_j(t) dt = 0 \quad \forall j \in \{1, \dots, J\}. \quad (44)$$

Alternatively, we can require orthogonality with respect to the sampled points (the resulting expansions are very similar for n large)

$$\frac{1}{n} \sum_{i=1}^n \Delta_J(i/n) \varphi_j(i/n) = 0 \quad \forall j \in \{1, 2, \dots, J\}. \quad (45)$$

With the remainder Δ_J , our regression model becomes

$$y = \mathbf{X}\theta_0 + \Delta_J + w, \quad (46)$$

where $\mathbf{X} \in \mathbb{R}^{n \times J}$, $w \sim \mathcal{N}(0, \sigma^2 I_n)$, and $\mathbf{X}^T \Delta_J = 0$. Recall, the LS estimator is given by

$$\begin{aligned} \hat{\theta} &= (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T y, \\ &= \theta_0 + (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T w. \end{aligned} \quad (47)$$

We can compute the prediction risk as follows

$$\begin{aligned} R_p(f) &= \frac{1}{n} \mathbb{E} \left\{ \sum_{i=1}^n \left(\hat{f}(i/n) - f(i/n) \right)^2 \right\}, \\ &= \frac{1}{n} \mathbb{E} \left\{ \|\mathbf{X} \hat{\theta} - \mathbf{X} \theta_0 - \Delta_J\|_2^2 \right\}, \\ &= \frac{1}{n} \mathbb{E} \left\{ \|\mathbf{X}(\theta - \hat{\theta}_0)\|_2^2 \right\} + \frac{1}{n} \|\Delta_J\|_2^2, \\ &= \frac{1}{n} \mathbb{E} \left\{ \|\mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T w\|_2^2 \right\} + \frac{1}{n} \|\Delta_J\|_2^2, \\ &= \frac{1}{n} \|\Delta_J\|_2^2 + \frac{\sigma^2}{n} \text{Tr}(\mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T). \end{aligned} \quad (48)$$

Finally, note that $\mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \in \mathbb{R}^{n \times n}$ is the orthogonal projector on the space spanned by the columns of \mathbf{X} . Hence its trace is always equal to J . This gives us the final form of the estimation risk at $f(\cdot)$, as a function of J ,

$$R_p(f) = \frac{\|\Delta_J\|_2^2}{n} + J \frac{\sigma^2}{n}. \quad (49)$$

In other words, the estimation risk associated with $f(\cdot)$ is a sum of two terms, both of which are dependent upon the choice of J :

- The first term is associated with the approximation error induced by the choice of the predictor $\{\varphi_j\}$. It is independent of the noise variance σ^2 , decreases with J .

We interpret it therefore a *bias* term.

- The second term depends on the noise level, and is related to the fluctuations that the noise induces in $\hat{\theta}$. It increases with the number of predictors J , as the fit becomes more unstable.

We interpret it therefore as a *variance* term.

Therefore, the optimal number of predictors to use, J^* is the one which minimizes the risk by striking a balance between bias and variance. In other words, we want to find the optimal point between under- and over-fitting the model. Note that how we choose the predictors themselves determines the rate at which the bias term goes to zero as J increases.

2.2.1. Example: The Fourier basis. In this example we select our set of predictors to be a Fourier basis

$$\varphi_j(t) = \sqrt{2} \cos((j-1)\pi t), \quad (50)$$




FIGURE 9. The effect of bias and variance on estimation risk.

for $t \in [0, 1]$. If $f(\cdot)$ is square-integrable, then it can be represented as an infinite series (converging in $L^2([0, 1])$, i.e. in mean square error)

$$f(t) = \sum_{j=1}^{\infty} \theta_{0j} \varphi_j(t). \quad (51)$$

However, if only J sinusoids are used, the remainder, Δ_J is

$$\Delta_J(t) = f(t) - \sum_{j=1}^J \theta_{0j} \varphi_j(t). \quad (52)$$

And, finally, the squared norm of Δ_J is, by orthogonality

$$\begin{aligned} \|\Delta_J\|_2^2 &= \sum_{i=1}^n \Delta_J(i/n)^2, \\ &\approx n \int_0^1 \Delta_J(t)^2 = n \int \sum_{j=J+1}^{\infty} \theta_{0j}^2. \end{aligned} \quad (53)$$

Here we replaced the sum by an integral, an approximation that is accurate for n large.

Note that the decay of the bias term mirrors the decay of the Fourier coefficients of f , by (53). In particular, if f is smooth, its Fourier coefficients decay faster, and hence the bias decays rapidly with J . In this case, the Fourier basis is a good set of features/predictors (a good dictionary) for our problem.

We now look at the case of Fourier predictors for a specific class of smooth functions, namely functions whose second derivative is square integrable. Formally we define

$$W(C) \equiv \left\{ \int_0^1 (f''(t))^2 dt \leq C^2 \right\}, \quad (54)$$

and we will consider estimation over $\Omega = W(C)$. This space is known in functional analysis as the ‘Sobolev ball of radius C and order 2.’




FIGURE 10. Depiction of the rate of decay of the bias term as a function of J .

In terms of Fourier coefficients, this set of smooth functions can be characterized as

$$\int_0^1 \left(\sum_{j=1}^{\infty} \pi^2 (j-1)^2 \theta_{0j} \varphi_j(t) \right)^2 dt \leq C \quad (55)$$

Or equivalently

$$\sum_{j=1}^{\infty} \pi^4 (j-1)^4 \theta_{0j}^2 \leq C. \quad (56)$$

Now, in order for this to happen, we must have (56) is satisfied,

$$\sum_{j=J+1}^{\infty} \theta_{0j}^2 \lesssim \frac{C'}{J^4}, \quad (57)$$

and hence, we can estimate a bound on the squared norm of the remainder term

$$\|\Delta_J\|_2^2 \approx n \sum_{j=J+1}^{\infty} \theta_{0j}^2 \lesssim \frac{nC'}{J^4}. \quad (58)$$


Therefore, the prediction risk for the set of functions $\Omega = W(C)$ is upper bounded

$$R_{p,*}(\Omega; \hat{\theta}) \lesssim \frac{C}{J^4} + \frac{\sigma^2 J}{n}. \quad (59)$$

The optimum value of J is achieved when the two terms are of the same order, or by setting to 0 the derivative with respect to J

$$\frac{\partial}{\partial J} \left\{ \frac{C}{J^4} + \frac{\sigma^2 J}{n} \right\} = \frac{\sigma^2}{n} - \frac{4C}{J^5}, \quad (60)$$

$$\therefore J_* \sim \left(\frac{n}{\sigma^2} \right)^{1/5}. \quad (61)$$


 FIGURE 11. Cutting off coefficients at J^* .

Finally, with the optimal choice of J , we obtain the upper bound for the prediction risk, in general,

$$R_{p,*}(\Omega; \hat{\theta}) \lesssim \left(\frac{\sigma^2}{n} \right)^{4/5}. \quad (62)$$

As in the standard parametric case, see (25), the risk depends on the ratio of the noise variance to the number of samples. However the decay with the number of samples is slower: $n^{-4/5}$ instead of n^{-1} . This is the price paid for not knowing in advance the p -dimensional space to which f belong. It can be proved that the exponent derived here is optimal.

2.3. Wavelet expansions. As we emphasized several times, the quality of our function estimation procedure is highly dependent on the choice of the features $\{\varphi_j\}$. More precisely, it depends on the ability to represent the signal of interest with a few elements of this dictionary. While the Fourier basis works well for smooth signals, it is not an adequate dictionary for many signals of interest. For instance, the Fourier expansion does not work very well for images.

Why is this the case? Let us reconsider what are the LS estimates for the Fourier coefficients. Using orthonormality of the Fourier basis, we have

$$\hat{\theta}_j^{\text{LS}} \approx \frac{1}{n} \sum_{i=1}^n \varphi_j(i/n) y_i = \theta_{0,j} + \hat{w}_j, \quad (63)$$

where $\hat{w}_i = n^{-1} \sum_{i=1}^n \varphi_j(i/n) y_i$. Figure 11 shows a cartoon of these coefficients.

In other words, each estimated coefficients is a sum of two contributions: the true Fourier coefficients $\theta_{0,j}$ and the noise contribution \hat{w}_j . Since the noise is white, its energy is equally spread across all Fourier modes. On the other hand, if the signal is smooth, its energy concentrates on low-frequency modes. By selecting a cut-off at J^* , we are sacrificing some true signal information, in order to get rid of most of

the noise. For frequencies higher than J^* , the noise energy surpasses any additional information these coefficients contain about the original signal we wish to estimate.

In other words, by selecting $J = J_*$, we are filtering out high frequencies in our measurements. In ‘time’ domain, this is essentially equivalent to averaging the observations over a sliding window of size of order J_*^{-1} . Formally, this is done by convolving the observations y , with some smooth kernel $K(\cdot)$,

$$\hat{f}(t) = \frac{1}{n} \sum_{i=1}^n K\left(\frac{i}{n} - t\right) y_i. \quad (64)$$

The details of the kernel $K(\cdot)$ can be worked out in detail, but what is important here is that it is $K(s)$ significantly different from 0 if and only if $|s| \leq J_*^{-1}$. This point of view gives a different perspective on the bias-variance tradeoff:

- For small J_* , we are averaging over large window, and hence reducing the variance of our estimates. On the other hand, we are introducing a large bias in favor of smooth signals.
- For large J_* we average over a large window. The estimate is less biased, but has a lot of variance.

For truly smooth signals, this approach to denoising is adequate. However, for many signals, the degree of smoothness changes dramatically from one point to the other of the signal. For instance, an image is mostly smooth, because of homogeneous surfaces corresponding to the same object or degree of illumination. However, they contain a lot of important discontinuities (e.g. edges) as well. Missing or smoothing out edges has a dramatic impact on the quality of reconstruction.

Smoothing with a kernel with uniform width produces a very bad reconstruction on such signals. If the width is large the image becomes blurred across these edges. If the width is small, it will not filter out noise efficiently in the smooth regions. A different predictor set must be used that adapts to different levels of smoothness in different point of the image.

Wavelets are one such basis. A wavelet expansion of a function allows for localization of frequency terms, which means high-frequency coefficients can be localized to edges, while smoother content of the image can be more concisely described with just a few low-frequency coefficients. Wavelets, as in our previous example of the Fourier basis, are an orthonormal basis of $[0, 1]$. The expansion is formed via two functions, the *father-wavelet*, or *scaling*, function, $\varphi(\cdot)$, and the *mother-wavelet* function, $\psi_{jk}(\cdot)$. The mother-wavelet function is used to generate set of self-similar functions which are composed of scaled and shifted versions of the mother-wavelet,

$$\psi_{jk}(t) = 2^{j/2} \psi(2^j t - k) \quad \text{where } j \in \{0, 1, 2, \dots\}, \quad (65)$$

$$k \in \{0, 1, \dots, 2^{j-1}\}.$$

Hence j is an index related to frequency, and k is related to position. The full wavelet expansion is then

$$f(t) = \theta_0 \varphi(t) + \sum_{j=0}^{\infty} \sum_{k=0}^{2^{j-1}} \theta_{0jk} \psi_{jk}(t). \quad (66)$$




FIGURE 12. Wavelet coefficients of a piecewise-continuous function for increasing scale levels.

There exist many families of wavelet functions, but the simplest among them is the Haar wavelet family. For the Haar wavelet, the wavelet functions are defined as

$$\varphi(t) = \begin{cases} 1 & \text{if } 0 \leq t < 1, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\psi(t) = \begin{cases} -1 & \text{if } 0 \leq t < \frac{1}{2}, \\ 1 & \text{if } \frac{1}{2} \leq t < 1, \\ 0 & \text{otherwise.} \end{cases}$$

In Fig. 12 we see an example wavelet expansion of a piecewise-continuous function. Larger magnitude wavelet coefficients will be located with the discontinuities in the original function across all scales.

Two problems arise naturally:

- (1) Unlike the Fourier basis, wavelet coefficients have no natural ordering of “importance”, since each wavelet coefficient describes the function at a certain length scale, and in a certain position. Hence, the simple idea of fitting all coefficients up to a certain maximum index J_* cannot be applied. If we select all

coefficients corresponding to all positions up to a certain maximum frequency, we will not exploit the spatial adaptivity property of the wavelet basis.

- (2) Any linear estimation procedure, that is also translation invariant can be represented as a convolution cf. (64), and thus incurs the problems outlined above. In order to treat differently edges from smooth regions in an image, a nonlinear procedure must be used.

The simplest approach that overcomes these problems is the wavelet denoising method that was developed in a sequence of seminal papers by David Donoho and Iain Johnstone [10, 11, 12, 13, 3]. The basic idea is to truncate, not according to the wavelet index, but according according to the magnitude of the measured wavelet coefficient. In the simplest implementation, we proceed in two steps. First we perform least squares estimation of each coefficient. In the case of orthogonal designs considered here, this yields

$$\tilde{y} = \frac{1}{n} \mathbf{X}^\top y = \theta + \tilde{w}, \quad (67)$$

Here $\tilde{w} = \mathbf{X}^\top w/n$ is again white noise $\tilde{w} \sim \mathcal{N}(0, (\sigma^2/n)\mathbf{I}_n)$. After this, coefficients are *thresholded*, independently,

$$\hat{\theta}_i = \begin{cases} \tilde{y}_i, & \text{if } |\tilde{y}_i| \geq \lambda, \\ 0, & \text{otherwise,} \end{cases}$$

The overall effect of this thresholding is to preserve large magnitude wavelet coefficients while zeroing those that are ‘below the noise level.’ Since larger coefficients corresponds to edged in the image, this approach seek to estimate higher frequencies near edges, only retaining low frequencies in smooth regions. This allows denoising without blurring across edges.

3. DENOISING WITH THRESHOLDING

In the last section we briefly described a denoising method, wavelet thresholding, that is can adapt to a degree of smoothness that varies across a signal (e.g. an image). In this section, we work out some basic properties of this method, under a simple signal model. Apart from being interesting per se, this analysis provides key insights for generalizing the same method to high-dimensional statistical estimation problems beyond denoising. For an in-depth treatment we refer, for instance, to [3, 10].

To recall the our set-up, we are considering the model

$$y = \mathbf{X}\theta + w, \quad (68)$$

where $y \in \mathbb{R}^n$, $\mathbf{X} \in \mathbb{R}^{n \times p}$ are observed, and we want to estimate the vector of coefficients $\theta \in \mathbb{R}^p$. The vector w is noise $w \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_n)$. We are focusing on orthogonal designs, i.e. on the case $n = p$ with $\mathbf{X}^\top \mathbf{X} = n, \mathbf{I}_{n \times n}$.

There is no loss of generality in carrying out least squares as a first step, which in this case reduces to

$$\tilde{y} = \frac{1}{n} \mathbf{X}^\top y = \theta + \tilde{w}, \quad \tilde{w} \sim \mathcal{N}(0, \frac{\sigma^2}{n} \mathbf{I}_{n \times n}). \quad (69)$$

In other words, in the case of orthogonal designs we can equivalently assume that the unknown object θ has been observed directly, with additive Gaussian noise.

Since we expect θ to be sparse, it is natural to return a sparse estimate $\widehat{\theta}$. In particular, if \tilde{y}_i is of the same order as the noise standard deviation σ , it is natural to guess that θ_i is actually very small or vanishing, and hence set $\widehat{\theta}_i = 0$. Two simple ways to implement this idea are ‘hard thresholding’ and ‘soft thresholding.’

Under *hard thresholding*, the estimate $\widehat{\theta} = (\widehat{\theta}_1, \dots, \widehat{\theta}_p)$ of θ is given by

$$\widehat{\theta}_i = \begin{cases} \tilde{y}_i & \text{if } |\tilde{y}_i| \geq \lambda, \\ 0 & \text{else.} \end{cases} \quad (70)$$

Under *soft-thresholding*, the estimate $\widehat{\theta}$ is given by

$$\widehat{\theta}_i = \begin{cases} \tilde{y}_i - \lambda & \text{if } \tilde{y}_i \geq \lambda, \\ 0 & \text{if } |\tilde{y}_i| \leq \lambda, \\ \tilde{y}_i + \lambda & \text{if } \tilde{y}_i \leq -\lambda. \end{cases} \quad (71)$$

These hard thresholding and soft thresholding functions are plotted in Fig. 13. While the two approaches have comparable properties (in particular, similar risk over sparse vectors), we shall focus here on soft thresholding since it is most easily generalizable to other estimation problems.




FIGURE 13. Soft thresholding (red) and hard thresholding (blue).

Note that both hard and soft thresholding depend on a threshold parameter that we denoted by λ . Entries below λ are set to zero: to achieve minimal risk, it is of course crucial to select an appropriate λ . Ideally, the threshold should cut-off the coefficients resulting from the noise, and hence we expect λ to be proportional to the noise standard deviation σ . In order to determine the optimal choice of λ , let us first consider the case $\theta = 0$. Note that, when $\theta = 0$, $\tilde{y} \sim \mathcal{N}(0, \frac{\sigma^2}{n} \mathbf{I}_n)$ is a vector with i.i.d. Gaussian entries. We claim that, in this case

$$\max_{i \in \{1, \dots, p\}} |\tilde{y}_i| \approx \sigma \sqrt{\frac{2 \log p}{n}}, \quad (72)$$

with probability very close to one¹

¹In this derivation we will be by choice somewhat imprecise, so as to increase readability. The reader is welcome to fill in the details, or to consult, for instance, [3, 10].

To see why this is the case, let $N(z) = \mathbb{E}\#\{i \in [p] : |\tilde{y}_i| \geq |z|\}$ be the expected number of coordinates in the vector \tilde{y} that are above level $|z|$, or below $-|z|$. By linearity of expectation, we have

$$N(z) = 2p \Phi\left(-\frac{n|z|}{\sigma}\right), \quad (73)$$

where $\Phi(x) = \int_{-\infty}^x e^{-t^2/2} dt / \sqrt{2\pi}$ is the Gaussian distribution function. Using the inequality $\Phi(-x) \leq e^{-x^2/2}/2$, valid for $x \geq 0$, we obtain $N(z) \leq p \exp(-nz^2/2\sigma^2)$. In particular, for any $\delta > 0$

$$\mathbb{P}\left\{\max_{i \in 1, \dots, p} |\tilde{y}_i| \geq \sigma \sqrt{\frac{2(1+\delta) \log p}{n}}\right\} \leq N\left(\sigma \sqrt{\frac{2(1+\delta) \log p}{n}}\right) \leq p^{-\delta}, \quad (74)$$

which vanishes as $p \rightarrow \infty$. Roughly speaking, this proves that $\max_{i \in 1, \dots, p} |\tilde{y}_i| \lesssim \sigma \sqrt{\frac{2 \log p}{n}}$ with high probability. A matching lower bound can be proved by a second moment argument and we leave it to the reader (or refer to the literature).




FIGURE 14. Sketch $\log N(z)$ (logarithm of the number of coordinates with noise level z).

Figure 14 reproduces the behavior of $\log N(z)$. The reader with a background in statistical physics has probably noticed the similarity between the present analysis and Derrida's treatment of the 'random-energy model' [14]. In fact the two models are identical and there is a close relationship between the problem addressed within statistical physics and estimation theory.

Figure 15 is a cartoon of the vector of observations \tilde{y} in the case in which the signal vanishes: $\theta = 0$. All the coordinates of \tilde{y} lie between $-\sigma \sqrt{\frac{2 \log p}{n}}$ and $+\sigma \sqrt{\frac{2 \log p}{n}}$. This suggests to set the threshold λ as to zero all the entries that are pure noise. This leads to the so-called following thresholding rule, proposed in [15]

$$\lambda = \sigma \sqrt{\frac{2 \log p}{n}}. \quad (75)$$




FIGURE 15. Cartoon of the vector of observations \tilde{y} when $\theta = 0$.

We now turn to evaluating the risk for such an estimator, when $\theta \neq 0$ is a sparse signal:

$$R(\theta; \hat{\theta}) = \mathbb{E}\{\|\theta - \hat{\theta}\|^2\} = \sum_{i=1}^p \mathbb{E}\{(\theta_i - \hat{\theta}_i)^2\}. \quad (76)$$

We can decompose this risk as

$$R = R_0 + R_{\neq 0},$$

where R_0 (respectively, $R_{\neq 0}$) is risk from entries θ_i that are zero (respectively, non-zero). The two contributions depend differently on λ : the contribution of zeros decreases with λ since for large λ more entries are set to 0. The contribution of non-zero entries instead increases with λ since large λ produces a larger bias, see Fig. 16 for a cartoon.




FIGURE 16. Risk $R_{\neq 0}$ in comparison to R_0 .

Under universal thresholding, since $\max_{i:\theta_i=0} |\tilde{y}_i| \lesssim \sigma \sqrt{\frac{2 \log p}{n}} = \lambda$, we have $R_0 \approx 0$. In order to evaluate the contribution of non-zero entries, we assume that θ is s_0 sparse, i.e., letting $\text{supp}(\theta) \equiv \{i \in [p] : \theta_i \neq 0\}$, we have $|\text{supp}(\theta)| \leq s_0$. Note that soft thresholding introduces a bias of size λ on these entries, as soon as they are sufficiently than λ . This gives an error per coordinate proportional to λ^2 (the variance contribution is negligible on these entries). This gives

$$R(\theta; \hat{\theta}) \approx R_{\neq 0} \approx s_0 \lambda^2 = \frac{s_0 \sigma^2}{n} (2 \log p). \quad (77)$$



FIGURE 17. This picture illustrates the universal thresholding. The solid dots represent $\tilde{\theta}$ and the bars represent \tilde{y} .

We can now step back and compare this result with the risk of least square estimation (25). Neglecting the factor $(2 \log p)$ which is small even for very high dimension, our formula for sparse vectors (77) is the same as for least squares, except that the dimension p is replaced by the number of non-zero entries s_0 . In other words, we basically achieve the same risk *as if* we knew a priori $\text{supp}(\theta)$ and run least squares on that support! The extra factor $(2 \log p)$ is the price we pay for not knowing where the support is. For sparse vectors, we achieve an impressive improvement over least squares.


Notice that this improvement is achieved simultaneously over all possible sparsity levels s_0 , and the estimator does not need to know a priori s_0 .

3.1. An equivalent analysis: Estimating a random scalar. There is a different, and essentially equivalent, way to analyze soft thresholding denoising. We will quickly sketch this approach because it provides an alternative point of view and, most importantly, because we will use some of its results in the next sections. We will omit spelling out the correspondence with the analysis in the last section.

We state this analysis in terms of a different –but essentially equivalent– problem. A source of information produces a random variable Θ in \mathbb{R} with distribution p_Θ , and we observe it corrupted by Gaussian noise. Namely, we observe Y given by

$$Y = \Theta + \tau Z, \quad (78)$$

where $Z \sim \mathcal{N}(0, 1)$ independent of Θ , and τ is the noise standard deviation. We want to estimate Θ from Y . A block diagram of this process is shown below. (A hint:



the correspondence with the problem in the previous section is obtained by setting $\tau = \sigma/\sqrt{n}$ and $p_\Theta = \frac{1}{p} \sum_{i=1}^p \delta_{\theta_i}$.)

We saw in the previous sections that sparse vectors can be used to model natural signals (e.g. images in wavelet domain). In the present framework, this can be modeled by the set of probability distributions that attribute mass at least $1 - \varepsilon$ to 0:

$$\mathcal{F}_\varepsilon = \{p_\Theta \in \mathcal{P} \mid p_\Theta(\{0\}) \geq 1 - \varepsilon\}, \quad (79)$$

where \mathcal{P} is the set of all probability distributions over the real line \mathbb{R} . Equivalently, \mathcal{F}_ε is the class of probability distributions that can be written as $p_\Theta = (1 - \varepsilon)\delta_0 + \varepsilon Q$ where δ_0 is the Dirac measure at 0 and Q is an arbitrary probability distribution.

The Bayes risk of an estimator $\hat{\Theta}$ is given by

$$R_B(p_\Theta; \hat{\Theta}) = \mathbb{E}\{[\hat{\Theta} - \Theta]^2\}. \quad (80)$$

In view of the interesting properties of soft thresholding, unveiled in the previous section, we will assume that Θ is obtained by soft thresholding Y . It is convenient at this point to introduce some notation for soft thresholding:

$$\eta(z; \lambda) = \begin{cases} z - \lambda & \text{if } z \geq \lambda, \\ 0 & \text{if } |z| \leq \lambda, \\ z + \lambda & \text{if } z \leq -\lambda. \end{cases} \quad (81)$$

With an abuse of notation, we write $R_B(p_\Theta; \lambda) = R_B(p_\Theta; \eta(\cdot; \lambda))$ for the Bayes risk of soft thresholding with threshold λ . Explicitly

$$R_B(p_\Theta; \lambda) = \mathbb{E}\{[\eta(Y; \lambda) - \Theta]^2\}. \quad (82)$$

We are interested in bounding the risk $R_B(p_\Theta; \lambda)$ for all sparse signals, i.e., in the present framework, for all the probability distributions $p_\Theta \in \mathcal{F}_\varepsilon$. We then consider the minimax risk:

$$R_*(\varepsilon; \tau^2) = \inf_\lambda \sup_{p_\Theta \in \mathcal{F}_\varepsilon} R_B(p_\Theta; \lambda). \quad (83)$$

First note that the class \mathcal{F}_ε is scale invariant. If $p_\Theta \in \mathcal{F}_\varepsilon$, also the probability distribution that is obtained by ‘stretching’ p_Θ by any positive factor s is in \mathcal{F}_ε . Hence the only scale in the problem is the noise variance τ^2 . It follows that

$$R_*(\varepsilon; \tau^2) = M(\varepsilon) \tau^2. \quad (84)$$

for some function $M(\varepsilon)$. Explicit formulae for $M(\varepsilon)$ can be found –for instance– in [16, Supplementary Material] or [17]. A sketch is shown in Fig. 18: in particular $M(\varepsilon) \approx 2\varepsilon \log \frac{1}{\varepsilon}$ as $\varepsilon \rightarrow 0$. By the same scaling argument as above, the optimal




FIGURE 18. Sketch of the minimax risk of soft thresholding $M(\varepsilon)$.

threshold λ takes the form

$$\lambda^* = \tau \ell(\varepsilon), \quad (85)$$

where the function $\ell(\varepsilon)$ can be computed as well and behaves as $\ell(\varepsilon) \approx \sqrt{2 \log \frac{1}{\varepsilon}}$ for small ε . Finally, the worst case signal distribution is

$$p_{\Theta}^* = (1 - \varepsilon)\delta_0 + \varepsilon\delta_{\infty}. \quad (86)$$

Note that the small ε behavior matches –as expected– the very sparse limit for vector denoising derived in the previous section. The correspondence is obtained by substituting $\varepsilon = s_0/p$ for the fraction of non-zero entries and noting that the vector risk is

$$R = pR_*(\varepsilon; \tau) = pM(\varepsilon)\tau^2. \quad (87)$$

When $s_0 \ll p$, we have $\varepsilon \approx 0$ and

$$R = pM(\varepsilon)\tau^2 \approx p 2\varepsilon \log \frac{1}{\varepsilon} \tau^2 \quad (88)$$

$$= p 2 \frac{s_0}{p} \log \frac{p}{s_0} \frac{\sigma^2}{n} \quad (89)$$

$$= \frac{2s_0\sigma^2}{n} \log \frac{p}{s_0}, \quad (90)$$

which matches the behavior derived earlier.

4. SPARSE REGRESSION

Up to now we have focused on estimating $\theta \in \mathbb{R}^p$ from observations of the form

$$y = \mathbf{X}\theta + w, \quad (91)$$

where $y \in \mathbb{R}^n$, $\mathbf{X} \in \mathbb{R}^{n \times p}$ are known and $w \in \mathbb{R}^n$ is an unknown noise vector. We focused in the previous case on orthogonal designs $n \geq p$ and $\mathbf{X}^T \mathbf{X} = n \mathbf{I}_{p \times p}$.

Over the last decade, there has been a lot of interest in the underdetermined case $n \ll p$, as general a \mathbf{X} as possible, which naturally emerges in many applications. It turns out that good estimation is possible provided θ is highly structured, and in particular when it is very sparse. Throughout, we let $\|\theta\|_0$ denote the ‘ ℓ_0 norm’ of θ , i.e. the number of non-zero entries in θ (note that this is not really a norm). The main outcome of the work in this area is that the number of measurements needs to scale with the number of non-zeros $\|\theta\|_0$ instead of the ambient dimension $p \gg \|\theta\|_0$. This setup is the so-called *sparse regression*, or *high-dimensional regression* problem.

4.1. Motivation. It is useful to overview a few scenarios where the above framework applies, and in particular the high-dimensional regime $n \ll p$ plays a crucial role.

Signal processing: An image can be modeled, for instance, by a function $f : [0, 1] \times [0, 1] \rightarrow \mathbb{R}$ if it is gray-scale. Color images requires three scalars at each point, three dimensional imaging requires to use a domain $[0, 1] \times [0, 1] \times [0, 1]$, and so on. Many imaging devices can be modeled (to a first order) as linear operators A collecting a vector $y \in \mathbb{R}^n$ corrupted by noise w denotes the noise.

$$y = Af + w, \quad (92)$$

To keep a useful example in mind, A can be the partial Fourier matrix, i.e. the operator computing a subset of Fourier coefficients. As emphasized in the

previous sections, the image f is often sparse in some domain, say wavelet transform. That is $f = T\theta$, where θ is sparse and T is the wavelet transform, or whatever sparsifying transform. This gives rise to the model

$$y = Af + w = (AT)\theta + w = \mathbf{X}\theta + w, \tag{93}$$

where $\mathbf{X} = AT$. Here n corresponds to the number of measurements, while p scales with the number of wavelet coefficient, and hence with the resolution that we want to achieve. The high-dimensional regime $n \ll p$ is therefore very useful as it corresponds to simpler measurements and higher resolution.

Machine learning: In web services, we often want to predict an unknown property of a user, on the basis of a large amount on known data about her. For instance, an online social network as Facebook, might want to estimate the income of its users, in order to display targeted advertisement. For each user i , we can construct a feature vector $x_i \in \mathbb{R}^p$, where e.g.,

$$x_i = (\text{age, location, number of friends, number of posts,} \\ \text{time of first post in a day, } \dots).$$

In a linear model, we assume

$$\underbrace{y_i}_{\text{income}} = \langle x_i, \theta \rangle + w_i, \tag{94}$$

Combining all users, we have

$$y = \mathbf{X}\theta + w, \tag{95}$$

where $y = (y_1, y_2, \dots, y_n)$ is a vector comprising all response variables (e.g. the customers' income), and \mathbf{X} is a matrix whose i -th row is the feature vector x_i of the i -th customer. Typically, one constructs feature vectors with tens of thousands of attributes, hence giving rise to $p = 10^4$ to 10^5 . On the other hand, in order to fit such a model, the response variable (income) y_i needs to be known for a set of users and this is often possible only for $n \ll p$ users.

Luckily, only a small subset of features is actually relevant to predict income, and hence we are led to use sparse estimation techniques.

4.2. The LASSO. The LASSO (Least Absolute Shrinkage and Selection Operator) presented in [18], also known as Basis Pursuit DeNoising (BPDN) [19, 20] is arguably the most successful method for sparse regression. The LASSO estimator is defined in terms of an optimization problem

$$\hat{\theta} = \underset{\theta \in \mathbb{R}^p}{\operatorname{argmin}} \left\{ \underbrace{\frac{1}{2n} \|y - \mathbf{X}\theta\|_2^2}_{\text{Residual sum of squares}} + \underbrace{\lambda \|\theta\|_1}_{\text{Regularizer}} \right\} \tag{LASSO}$$

The term $\mathcal{L}(\theta) = \frac{1}{2n} \|y - \mathbf{X}\theta\|_2^2$ is the ordinary least squares cost function, and the regularizer $\lambda \|\theta\|_1$ promotes sparse vectors by penalizing coefficients different from 0. Note that the optimization problem is convex and hence it can be solved efficiently: we wil discuss a simple algorithm in the following.

To gain insight as to why the LASSO is well-suited for sparse regression, let us start by revisiting the case of orthogonal designs, namely $n \geq p$ and

$$\mathbf{X}^\top \mathbf{X} = n \mathbf{I}_{n \times n}. \quad (96)$$

Rewriting $\mathcal{L}(\theta)$:

$$\begin{aligned} \mathcal{L}(\theta) &= \frac{1}{2n} \langle y - \mathbf{X}\theta, y - \mathbf{X}\theta \rangle, \\ &= \frac{1}{2n} \left\langle y - \mathbf{X}\theta, \frac{1}{n} \mathbf{X}\mathbf{X}^\top (y - \mathbf{X}\theta) \right\rangle, \\ &= \frac{1}{2n^2} \langle \mathbf{X}^\top y - n\theta, \mathbf{X}^\top y - n\theta \rangle, \\ &= \frac{1}{2} \left\| \theta - \frac{1}{n} \mathbf{X}^\top y \right\|^2, \\ &= \frac{1}{2} \|\theta - \tilde{y}\|^2, \quad \text{where } \tilde{y} = \frac{1}{n} \mathbf{X}^\top y. \end{aligned} \quad (97)$$

Thus, in this case, the LASSO problem is equivalent to

$$\text{minimize } \sum_{i=1}^p \left\{ \frac{1}{2} |\tilde{y}_i - \theta_i|^2 + \lambda |\theta_i| \right\}. \quad (98)$$

This is a ‘separable’ cost function, and we can minimize each coordinate separately. Let $F(\theta_i) = \frac{1}{2}(\tilde{y} - \theta_i)^2 + \lambda|\theta_i|$. Now,

$$\frac{\partial F}{\partial \theta_i} = \theta_i - \tilde{y}_i + \lambda \text{sign}(\theta_i), \quad (99)$$

where $\text{sign}(\cdot)$ denotes the sign function shown in Figure 19. Note that $|\theta_i|$ is non-





FIGURE 19. The sign function.

differentiable at $\theta_i = 0$. How should we interpret its derivative $\text{sign}(\theta_i)$ in this case? For convex functions (which is the case here) the derivative can be safely replaced by the ‘subdifferential,’ i.e. the set of all possible slopes of tangent lines at θ_i that stay below the graph of the function to be differentiated. The subdifferential coincides with the usual derivative where the function is differentiable. For the function $\theta_i \rightarrow |\theta_i|$, it is an easy exercise to check that the subdifferential at $\theta_i = 0$ is given by the interval $[-1, 1]$. In other words, we can think of Figure 19 as the correct graph of the subdifferential of $|\theta_i|$ if we interpret its value at 0 as given by the whole interval $[-1, 1]$.

The minimizer of $F(\theta_i)$ must satisfy

$$\tilde{y}_i = \theta_i + \lambda \operatorname{sign}(\theta_i). \quad (100)$$

Hence we can obtain the minimizer as a function of \tilde{y}_i by adding θ_i to the graph in figure (19) and flipping the axis. The result is plotted in the next figure. The



reader will recognize the soft thresholding function $\eta(\cdot; \lambda)$ already encountered in the previous section. Summarizing, in the case of orthogonal designs, the LASSO estimator admits the explicit representation

$$\hat{\theta} = \eta\left(\frac{1}{n}\mathbf{X}^\top y; \lambda\right), \quad (101)$$

where it is implicitly understood that the soft thresholding function is applied component-wise to the vector $(1/n)\mathbf{X}^\top y$. As we saw in the previous section, component-wise soft-thresholding has nearly optimal performances on this problem, and hence the same holds for the LASSO.

In the high-dimensional setting $p \gg n$ and \mathbf{X} is obviously not orthogonal, and the LASSO estimator is non-explicit. Nevertheless it can be computed efficiently, and we will discuss next a simple algorithm that is guaranteed to converge. It is an example of a generic method for convex optimization known as a ‘subgradient’ or ‘projected gradient’ approach. The important advantage of these algorithms (and more generally of ‘first order methods’) is that their complexity per iteration scales only linearly in the dimensions p of the problem, and are hence well suited for high-dimensional applications [21]. They are not as fast to converge as –for instance– Newton’s method, but this is often not crucial. For statistical problems a ‘low precision’ solution is often as good as a ‘high precision’ one since in any case there is an unavoidable statistical error to deal with.

We want to minimize the cost function

$$F(\theta) = \frac{1}{2n}\|y - \mathbf{X}\theta\|^2 + \lambda\|\theta\|_1. \quad (102)$$

At each iteration, the algorithm constructs an approximation $\theta^{(t)}$ of the minimizer $\hat{\theta}$. In order to update this state, the idea is to construct an upper bound to $F(\theta)$ that is

easy to minimize and is a good approximation of $F(\theta)$ close to $\theta^{(t)}$. Rewriting $\mathcal{L}(\theta)$:

$$\begin{aligned}\mathcal{L}(\theta) &= \frac{1}{2n} \|y - \mathbf{X}\theta\|_2^2, \\ &= \frac{1}{2n} \|y - \mathbf{X}\theta^{(t)} - \mathbf{X}(\theta - \theta^{(t)})\|_2^2, \\ &= \frac{1}{2n} \|y - \mathbf{X}\theta^{(t)}\|_2^2 - \frac{1}{n} \langle \mathbf{X}(\theta - \theta^{(t)}), y - \mathbf{X}\theta^{(t)} \rangle + \frac{1}{2n} \|\mathbf{X}(\theta - \theta^{(t)})\|_2^2.\end{aligned}\quad (103)$$

Note that the first two terms are ‘simple’ in that they are linear in θ . The last term is ‘small’ for θ close to $\theta^{(t)}$ (quadratic in $(\theta - \theta^{(t)})$). We will upper bound the last term. Suppose the largest eigenvalue of $\frac{1}{n}\mathbf{X}^\top\mathbf{X}$ is bounded by L ,

$$\lambda_{\max}\left(\frac{1}{n}\mathbf{X}^\top\mathbf{X}\right) \leq L, \quad (104)$$

and let $v = \frac{1}{n}\mathbf{X}^\top(y - \mathbf{X}\theta^{(0)})$. Then

$$\begin{aligned}\mathcal{L}(\theta) &= \frac{1}{2n} \|y - \mathbf{X}\theta^{(t)}\|_2^2 - \langle v, \theta - \theta^{(t)} \rangle + \frac{1}{2} \left\langle \theta - \theta^{(t)}, \frac{1}{n}\mathbf{X}^\top\mathbf{X}(\theta - \theta^{(t)}) \right\rangle, \\ &\leq \frac{1}{2n} \|y - \mathbf{X}\theta^{(t)}\|_2^2 - \langle v, \theta - \theta^{(t)} \rangle + \frac{L}{2} \|\theta - \theta^{(t)}\|_2^2, \\ &= \underbrace{\frac{1}{2n} \|y - \mathbf{X}\theta^{(t)}\|_2^2 - \frac{1}{2L} \|v\|_2^2}_{\triangleq C} + \frac{L}{2} \|\theta - \theta^{(t)} - \frac{1}{L}v\|_2^2.\end{aligned}\quad (105)$$

We therefore obtain the following upper bound, whereby C is a constant independent of θ ,

$$F(\theta) \leq C + \lambda \|\theta\|_1 + \frac{L}{2} \left\| \theta - \theta^{(0)} - \frac{v}{L} \right\|_2^2. \quad (106)$$

We compute the next iterate $\theta^{(t+1)}$ by minimizing the above upper bound:

$$\begin{aligned}\text{minimize} \quad & \frac{\lambda}{L} \|\theta\|_1 + \frac{1}{2} \|\theta - \tilde{\theta}^{(0)}\|_2^2, \\ & \tilde{\theta}^{(0)} = \theta^{(0)} + \frac{1}{nL} \mathbf{X}^\top(y - \mathbf{X}\theta^{(0)}).\end{aligned}\quad (107)$$

We already solved this problem when discussing the case of orthogonal designs. The solution is given by the soft thresholding operator:

$$\theta^{(t+1)} = \eta \left(\theta^{(t)} + \frac{1}{nL} \mathbf{X}^\top(y - \mathbf{X}\theta^{(t)}); \frac{\lambda}{L} \right). \quad (108)$$

This yields an iterative procedure known as Iterative Soft Thresholding that can be initialized arbitrarily, e.g. with $\theta^{(0)} = 0$. This algorithm is guaranteed to always converge, as shown in [22, 23], in the sense that

$$F(\theta^{(t)}) - F(\hat{\theta}) \leq \frac{\text{constant}}{t}. \quad (109)$$

Note that this is much slower than the rate achieved by Newton’s method. However it can be proved that no first order method (i.e. no method using only gradient information) can achieve global convergence rate faster than $1/t^2$ for any problem

in the class of the LASSO. We refer to [21] for a recent introduction to first order methods.




FIGURE 20. Subgradient approach: construct a convenient upper bound to $F(\theta)$.

To conclude this section, it is instructive to quickly consider two special cases: $\lambda \rightarrow \infty$ and $\lambda \rightarrow 0$. We rewrite the minimization problem as

$$\begin{aligned} \hat{\theta} &= \underset{\theta \in \mathbb{R}^p}{\operatorname{argmin}} F(\theta), \\ \frac{1}{\lambda} F(\theta) &= \frac{1}{2n\lambda} \|y - \mathbf{X}\theta\|_2^2 + \|\theta\|_1. \end{aligned} \quad (110)$$

When $\lambda \rightarrow \infty$, the first term vanishes and $\hat{\theta} \rightarrow \operatorname{argmin}_{\theta} \|\theta\|_1 = 0$. In fact, $\hat{\theta} = 0$ for all $\lambda \geq \lambda_*$ for some critical λ_* . When $\lambda \rightarrow 0$, the weight in front of the first term goes to infinity, and hence the equality $y = \mathbf{X}\theta$ is enforced strictly. In the high-dimensional regime $p > n$, this linear system is underdetermined and has multiple solution. The most relevant is selected by minimizing the ℓ_1 norm. In other words, as $\lambda \rightarrow 0$, the LASSO estimator $\hat{\theta}$ converges to the solution of the following problem (known as ‘basis pursuit’)

$$\begin{aligned} &\text{minimize} && \|\theta\|_1, \\ &\text{subject to} && y = \mathbf{X}\theta. \end{aligned} \quad (111)$$

4.3. Behavior of the LASSO under Restricted Isometry Property. A significant amount of theory has been developed to understand and generalize the remarkable properties of the LASSO estimator and its empirical success. The theory establishes certain optimality properties under suitable assumptions on the design matrix \mathbf{X} . The most popular of these assumptions goes under the name of restricted isometry property (RIP) and was introduced in the groundbreaking work of Candés, Tao and collaborators [24, 25]. Several refinements of this condition were developed in recent years (the restricted eigenvalue condition of [26], the compatibility condition of [5] and so on).

In order to motivate the RIP, we notice that the LASSO estimator performs well when the columns of the matrix \mathbf{X} are orthogonal, $\mathbf{X}^T \mathbf{X} = n \mathbf{I}_{n \times n}$. Indeed this is the case of orthogonal designs explored above. The orthogonality condition $\mathbf{X}^T \mathbf{X} = n \mathbf{I}_{n \times n}$ is equivalent to

$$\|\mathbf{X}v\|^2 = n\|v\|^2 \quad \text{for all } v \in \mathbb{R}^p \quad (112)$$

This is of course impossible in the high-dimensional regime $p > n$ (indeed the null space of \mathbf{X} has dimension at least $p - n$). The idea is to relax this condition, by requiring that \mathbf{X} is “almost orthogonal” instead of orthogonal, and only when it acts on sparse vectors. Explicitly, we say that \mathbf{X} satisfies the condition $\text{RIP}(k, \delta)$ for some integer k and $\delta \in (0, 1)$ if

$$(1 - \delta)\|v\|^2 \leq \frac{1}{n}\|\mathbf{X}v\|^2 \leq (1 + \delta)\|v\|^2 \quad \text{for all } v \in \mathbb{R}^p \text{ with } \|v\|_0 \leq k \quad (\text{RIP property}) \quad (113)$$

It is possible to show that this definition is non-empty and indeed—in a certain sense—most matrices satisfy it. For instance if \mathbf{X} has iid entries $\mathbf{X}_{ij} \sim \text{Unif}\{+1, -1\}$ or $\mathbf{X}_{ij} \sim \mathcal{N}(0, 1)$, then with high probability, \mathbf{X} satisfies $\text{RIP}(k, \delta)$ for a fixed δ and $n \geq Ck \log \frac{p}{k}$. The RIP property has been established for a large number of matrix constructions. For instance partial Fourier matrices² satisfy RIP with high probability for $n \geq C_1 k (\log p)^4$, as shown in [27].

The following theorem illustrates the utility of RIP matrices for sparse estimation. It is a simplified version of stronger results established in [26] (without any attempt at reproducing optimal constants, or the explicit dependence of all the quantities). Results of the same nature were proved earlier in [25] for a closely related estimator, known as the ‘Dantzig selector.’

Theorem 2 (Candes, Tao 2006 and Bickel, Ritov, Tsybakov 2009). *If θ is s_0 -sparse and \mathbf{X} satisfies $\text{RIP}(10s_0, 0.1)$, then, by choosing $\lambda = \sigma \sqrt{\frac{5 \log p}{n}}$, we have, with high probability for a suitable constant $C > 0$,*

$$\|\hat{\theta} - \theta\|_2^2 \leq \frac{C s_0 \sigma^2}{n} \log p. \quad (114)$$

A few observations are in order. It is—once again—instructive to compare this bound with the risk of ordinary least squares, cf. Eq. (25). Apart from the $\log p$ factor, the error scales as if θ_0 was s_0 -dimensional. As in the case of orthogonal designs discussed above, we obtain roughly the same scaling as if the support of θ was known. Also the choice of λ scales as in the case where \mathbf{X} is orthogonal.

Finally, as $\sigma \rightarrow 0$, we have $\hat{\theta} \rightarrow \theta$ provided the RIP condition is satisfied. As mentioned above, this happens for random design matrices if $n \geq C s_0 \log p$. In other words, we can reconstruct exactly an s_0 -sparse vector from about $s_0 \log p$ random linear observations.

4.3.1. Modeling the design matrix \mathbf{X} . The restricted isometry property and its refinements/generalizations allow to build a develop a powerful theory of high-dimensional statistical estimation (both in the context of linear regression and beyond). This approach has a number of strengths:

- (a) Given a matrix \mathbf{X} , we can characterize it in terms of its RIP constant, and hence obtain a bound on the resulting estimation error. The bound holds uniformly over all signals θ .
- (b) The resulting bound is often nearly optimal.

²That is, matrices obtained by subsampling randomly the rows of the $p \times p$ discrete Fourier transform.

- (c) Many class of random matrices of interest have been proved to possess RIP.
- (d) RIP allows to decouple the analysis of the statistical error, e.g. the risk of the LASSO estimator $\hat{\theta}$, (which is the main object interest of statisticians) from the development of algorithms to compute $\hat{\theta}$ (which is the focus within the optimization community).

The RIP theory has also some weaknesses. It is useful to understand them since this exercise leads to several interesting research directions that are –to a large extent– still open:

- (a) In practice, given a matrix \mathbf{X} it is NP-hard to whether it has RIP. Hence, one has often to rely on the intuition provided by random matrix constructions.
- (b) The resulting bounds typically optimal *within a constant*, that can be quite large. This makes it difficult to compare different estimators for the same problem. If estimator $\hat{\theta}^{(1)}$ has risk that is –say– twice as large as the one of $\hat{\theta}^{(2)}$, this is often not captured by this theory.
- (c) As a special case of the last point, RIP theory provides little guidance for the practically important problem of selecting the right amount of regularization λ . It is observed in practice that changing λ by a modest amount has important effects on the quality of estimation, but this is hardly captured by RIP theory.
- (d) Since RIP theory aims at bounding the risk uniformly over all (sparse) vectors θ , it is typically driven by the ‘worst case’ vectors, and is overly conservative for most θ ’s.

Complementary information on the LASSO, and other high-dimensional estimation methods, can be gathered by studying simple random models for the design matrix \mathbf{X} . This will be the object of the next lecture.

5. RANDOM DESIGNS AND APPROXIMATE MESSAGE PASSING

In this lecture we will revisit the linear model (91) and the LASSO estimator, while assuming a very simple probabilistic model for the design matrix \mathbf{X} . Before proceeding, we should therefore ask: Is there any application for which probabilistic design matrices are well suited? Two type of examples come to mind

- In statistics and machine learning, we are given pairs (response variable, covariate vector), $(y_1, x_1), \dots, (y_n, x_n)$ and postulate a relationship as for instance in Eq. (94). These pairs can often be thought as samples from a larger ‘population,’ e.g. customers of a e-commerce site are samples of a population of potential customers.

One way to model this, is to assume that the covariate vectors x_i ’s, i.e. the rows of \mathbf{X} are i.i.d. samples from a distribution.

- In compressed sensing, the matrix \mathbf{X} models a sensing or sampling device, that is designed within some physical constraints. Probabilistic constructions have been proposed and implemented by several authors, see e.g. [28] for an example. A cartoon example of these constructions is obtained by sampling i.i.d. random rows from the discrete $p \times p$ Fourier transform.

In other words, random design matrices \mathbf{X} with i.i.d. rows can be used to model several applications. Most of the work has however focused on the special case in

which the rows are i.i.d. with distribution $\mathcal{N}(0, \mathbf{I}_{p \times p})$. Equivalently, the matrix \mathbf{X} has i.i.d. entries $\mathbf{X}_{i,j} \sim \mathcal{N}(0, 1)$. Despite its simplicity, this model has been an important playground for the development of many ideas in compressed sensing, starting with the pioneering work of Donoho [29], and Donoho and Tanner [30, 31]. Recent years have witnessed an explosion of contributions also thanks to the convergence of powerful ideas from high-dimensional convex geometry and Gaussian processes, see e.g. [32, 33, 34, 35, 36]. Non-rigorous ideas from statistical physics were also used in [37, 38, 39, 40].

Here we follow a rigorous approach that builds upon ideas from statistical physics, information theory and graphical models, and is based on the analysis of an highly efficient reconstruction algorithm. We will sketch the main ideas referring to [16] for the original idea, to [17, 41, 42] for the analysis of the LASSO, and to [7, 43, 44] for extensions. This approach was also used in [45] to establish universality of the compressed sensing phase transition for non-Gaussian i.i.d. entries $\mathbf{X}_{i,j}$.

5.1. Message Passing algorithms. The plan of our analysis is as follows:

- (1) We define an approximate message passing (AMP) algorithm to solve the LASSO optimization problem. The derivation presented here starts from the subgradient method described in Section 4.2 and obtain a slight –but crucial– modification of the same algorithm. Also in this case the algorithm is iterative and computes a sequence of iterates $\{\theta^{(t)}\}$.

An alternative approach (susceptible of generalizations –for instance– to Bayesian estimation) is presented in [46].

- (2) Derive an exact asymptotic characterization of the same algorithm as $n, p \rightarrow \infty$, for t fixed. The characterization is given in terms of the so-called state evolution method developed rigorously in [41] (with generalizations in [43, 45]).
- (3) Prove that AMP converges fast to the optimized $\hat{\theta}$, namely with high probability as $n, p \rightarrow \infty$ we have $\|\theta^{(t)} - \hat{\theta}\|_2^2/p \leq c_1 e^{-c_2 t}$, with c_1, c_2 two dimension-independent constants. A full proof of this step can be found in [42].
- (4) Select t a large enough constant and use the last two result to deduce properties of the optimizer $\hat{\theta}$.

We next provide a sketch of the above steps. We start by considering iterative soft thresholding with $L = 1$:

$$\begin{cases} \theta^{(t+1)} = \eta(\theta^{(t)} + \frac{1}{n} \mathbf{X}^\top r^{(t)}; \gamma_t), \\ r^{(t)} = y - \mathbf{X} \theta^{(t)}, \end{cases} \quad (115)$$

where we introduced the additional freedom of an iteration-dependent threshold γ_t (instead of λ). Component-wise, the iteration takes the form

$$\begin{cases} \theta_i^{(t+1)} = \eta(\theta_i^{(t)} + \frac{1}{n} \sum_{a=1}^n \mathbf{X}_{ai} r_a^{(t)}; \gamma_t), \\ r_a^{(t)} = y_a - \sum_{i=1}^p \mathbf{X}_{ai} \theta_i^{(t)}. \end{cases} \quad (116)$$

We next derive a message passing version of this iteration³ (we refer for instance to [47, 48] for background). The motivation for this modification is that message

³We use the expression ‘message passing’ in the same sense attributed in information theory and graphical models.

passing algorithms have appealing statistical properties. For instance, they admit an exact asymptotic analysis on locally tree-like graphs. While –in the present case– the underlying graph structure is not locally tree-like, the conclusion (exact asymptotic characterization) continues to hold.

In order to define the message-passing version, we need to associate a factor graph to the LASSO cost function:

$$F(\theta) = \frac{1}{2n} \sum_{a=1}^n (y_a - \langle x_a, \theta \rangle)^2 + \lambda \sum_{i=1}^p |\theta_i|. \quad (117)$$

Following a general prescription from [48], we associate a factor node to each term $(y_a - \langle x_a, \theta \rangle)^2 / (2n)$ in the cost function indexed by $a \in \{1, 2, \dots, n\}$ (we do not need to represent the singletons $|\theta_i|$ by factor nodes), and we associate a variable node to each variable, indexed by $i \in \{1, 2, \dots, p\}$. We connect factor node a and variable node i by an edge (a, i) if and only if term a depends on variable θ_i , i.e. if $\mathbf{X}_{ai} \neq 0$. Note for Gaussian design matrices, all the entries \mathbf{X}_{ai} are non-zero with probability one. Hence, the resulting factor graph is a complete bipartite graph with n factor nodes and p variable nodes.

The message-passing version of the iteration (116) has iteration variables (messages) associated to directed edges of the factor graph. Namely, for each edge (a, i) we introduce a message $r_{a \rightarrow i}^{(t)}$ and a message $\theta_{i \rightarrow a}^{(t)}$. We replace the update rule (116) by the following

$$\begin{cases} \theta_{i \rightarrow a}^{(t+1)} = \eta \left(\frac{1}{n} \sum_{b \in [n] \setminus a} \mathbf{X}_{bi} r_{b \rightarrow i}^{(t)}; \gamma_t \right), \\ r_{a \rightarrow i}^{(t)} = y_a - \sum_{j \in [p] \setminus i} \mathbf{X}_{aj} \theta_{j \rightarrow a}^{(t)}. \end{cases} \quad (118)$$

The key property of this iteration is that an outgoing message from node α is updated by evaluating a function of all messages incoming in the same node α , except the one along the same edge. An alternative derivation of this iteration follows by considering the standard belief propagation algorithm (in its sum-product or min-sum forms), and using a second order approximation of the messages as in [46].

Note that, with respect to standard iterative soft thresholding, cf. Eq. (115), the algorithm (118) has higher complexity, since it requires to keep track of $2np$ messages, as opposed to the $n + p$ variables in Eq. (115). Also, there is obvious interpretation to the fixed points of the iteration (118).

It turns out that a simpler algorithm can be defined, whose state as dimension $n + p$ as for iterative soft thresholding, but tracks closely the iteration (118). This builds on the remark that the messages $\theta_{i \rightarrow a}^{(t)}$ issued from a node i do not differ to much, since their definition in Eq. (118) only differ in one out of n terms. A similar argument applies to the messages $r_{a \rightarrow i}^{(t)}$ issued by node a . We then write $\theta_{i \rightarrow a}^{(t)} = \theta_i^{(t)} + \delta \theta_{i \rightarrow a}^{(t)}$, $r_{a \rightarrow i}^{(t)} = r_a^{(t)} + \delta r_{a \rightarrow i}^{(t)}$ and linearize the iteration (118) in $\{\delta \theta_{i \rightarrow a}^{(t)}\}$, $\{\delta r_{a \rightarrow i}^{(t)}\}$. After eliminating these quantities [46], the resulting iteration takes the form, known as approximate message passing (AMP)

$$\begin{cases} \theta^{(t+1)} = \eta \left(\theta^{(t)} + \frac{1}{n} \mathbf{X}^T r^{(t)}; \gamma_t \right), \\ r^{(t)} = y - \mathbf{X} \theta^{(t)} + \mathbf{b}_t r^{(t-1)}, \end{cases} \quad (\text{AMP})$$

where $\mathbf{b}_t \equiv \|\theta^{(t)}\|_0/n$ is a scalar. In other words we recovered iterative soft thresholding except for the memory term $\mathbf{b}_t r^{(t-1)}$ that is straightforward to evaluate. In the context of statistical physics, a similar correction is known as the Onsager term. Remarkably, this memory term changes the statistical behavior of the algorithm.

It is an instructive exercise (left to the reader) to prove that fixed points of the AMP algorithm (with $\gamma_t = \gamma_*$ fixed) are minimizers of the LASSO. In particular, for Gaussian sensing matrices, such minimizer is unique with probability one.

We notice in passing that there is nothing special about the least squares objective, or the ℓ_1 regularization in our derivation. Indeed similar ideas were developed and applied to a large number of problems, see [49, 49, 50, 7, 51, 52, 53, 54] for a very incomplete list of examples.

5.2. Analysis of AMP and the LASSO. We next carry out a heuristic analysis of AMP, referring to [41] for a rigorous treatment that uses ideas developed by Bolthausen in the context of mean-field spin glasses [55].

We use the message passing version of the algorithm, cf. Eq. (118) and we will use the assumption that the pairs $\{(r_{a \rightarrow i}^{(t)}, \mathbf{X}_{ai})\}_{a \in [n]}$ are “as if” independent, and likewise for $\{(\theta_{a \rightarrow i}^{(t)}, \mathbf{X}_{ai})\}_{i \in [p]}$. This assumption is only approximately correct, but leads to the right asymptotic conclusions.

Consider the first equation in (118), and further assume (this assumption will be verified inductively)

$$\mathbb{E}(r_{a \rightarrow i}^{(t)}) = \mathbf{X}_{ai} \theta_i, \quad \text{Var}(r_{a \rightarrow i}^{(t)}) = \tau_t^2. \quad (119)$$

Letting $\tilde{r}_{a \rightarrow i}^{(t)} \equiv r_{a \rightarrow i}^{(t)} - \mathbb{E}(r_{a \rightarrow i}^{(t)})$, the argument of $\eta(\cdot; \gamma_t)$ in Eq. (118) can be written as

$$\frac{1}{n} \sum_{b \in [n] \setminus a} \mathbf{X}_{bi} r_{b \rightarrow i}^{(t)} = \frac{1}{n} \sum_{b \in [n] \setminus a} \mathbf{X}_{bi}^2 \theta_i + \frac{1}{n} \sum_{b \in [n] \setminus a} \mathbf{X}_{bi} \tilde{r}_{b \rightarrow i}^{(t)} \approx \theta_i + \frac{\tau_t}{\sqrt{n}} Z_{i \rightarrow a}^{(t)}, \quad (120)$$

where, by central limit theorem, $Z_{i \rightarrow a}^{(t)}$ is approximately distributed as $\mathcal{N}(0, 1)$.

Rewriting the first equation in (118), we obtain

$$\theta_{i \rightarrow a}^{(t+1)} = \eta\left(\theta_i + \frac{\tau_t}{\sqrt{n}} Z_{i \rightarrow a}^{(t)}; \gamma_t\right). \quad (121)$$

In the second message equation, we substitute $y_a = w_a + \sum_{j=1}^p \mathbf{X}_{aj} \theta_j$, thus obtaining

$$r_{a \rightarrow i}^{(t+1)} = w_a + \mathbf{X}_{ai} \theta_i - \sum_{j \in [p] \setminus i} \mathbf{X}_{aj} (\theta_{j \rightarrow a}^{(t+1)} - \theta_j). \quad (122)$$

The first and the last terms have 0 mean thus confirming the induction hypothesis $\mathbb{E}(r_{a \rightarrow i}^{(t+1)}) = \mathbf{X}_{ai} \theta_i$. The variance of $r_{a \rightarrow i}^{(t+1)}$ is given by (neglecting sublinear terms)

$$\tau_{t+1}^2 = \sigma^2 + \sum_{j=1}^p \left[\eta\left(\theta_j + \frac{\tau_t}{\sqrt{n}} Z_j; \gamma_t\right) - \theta_j \right]^2. \quad (123)$$

It is more convenient to work with the rescaled quantities $\tilde{\theta}_i = \theta_i \sqrt{n}$ and $\tilde{\gamma}_t = \gamma_t \sqrt{n}$ (this allows us to focus on the most interesting regime, whereby θ_i is of the same order

as the noise level τ_t/\sqrt{n}). Using the scaling property of the thresholding function $\eta(ax, a\gamma) = a\eta(x, \gamma)$, the last equation becomes

$$\tau_{t+1}^2 = \sigma^2 + \frac{1}{n} \sum_{j=1}^p \left[\eta(\tilde{\theta}_j + \tau_t Z_j; \tilde{\gamma}_t) - \tilde{\theta}_j \right]^2. \quad (124)$$

We now define the probability measure p_Θ as the asymptotic empirical distribution of $\tilde{\theta}$, $p^{-1} \sum_{j=1}^p \delta_{\tilde{\theta}_j}$ (formally, we assume that $p^{-1} \sum_{j=1}^p \delta_{\tilde{\theta}_j}$ converges weakly to p_Θ , and that low order moments converge as well). We also let $\delta = \lim_{n \rightarrow \infty} (n/p)$ be the asymptotic aspect ratio of \mathbf{X} . We then obtain

$$\tau_{t+1}^2 = \sigma^2 + \frac{1}{\delta} \mathbb{E} \left\{ (\eta(\Theta + \tau_t Z; \tilde{\gamma}_t) - \Theta)^2 \right\}, \quad (125)$$

where expectation is with respect to $\Theta \sim p_\Theta$ independent of $Z \sim \mathcal{N}(0, 1)$. The last equation is known as *state evolution*: despite the many unjustified assumptions in our derivation, it can be proved to correctly describe the $n, p \rightarrow \infty$ asymptotics of the message passing algorithm (118) as well as of the AMP algorithm.

Reconsidering the above derivation, we can derive asymptotically exact expressions for the risk at θ for of the AMP estimator $\theta^{(t+1)}$. Namely, we define the asymptotic risk

$$R_\infty(\theta; \theta^{(t+1)}) = \lim_{n, p \rightarrow \infty} \mathbb{E} \left\{ \|\theta - \theta^{(t)}\|^2 \right\}, \quad (126)$$

the limit being taken along sequences of vectors θ with converging empirical distribution. Then we claim that the limit exists and is given by

$$R_\infty(\theta; \theta^{(t+1)}) = \frac{1}{\delta} \mathbb{E} \left\{ (\eta(\Theta + \tau_t Z; \tilde{\gamma}_t) - \Theta)^2 \right\}, \quad (127)$$

or, equivalently,

$$R_\infty(\theta; \theta^{(t+1)}) = \tau_{t+1}^2 - \sigma^2. \quad (128)$$

Thus, apart from an additive constant, τ_t^2 coincides with risk and the latter can be tracked using state evolution.


In [42], it is proved that the AMP iterates $\theta^{(t)}$ converge rapidly to the LASSO estimator $\hat{\theta}$. We are therefore led to consider the large t behavior of τ_t , which yields the risk of the LASSO, or –equivalently– the risk of AMP after a sufficiently large (constant in n, p) number of iterations. Before addressing this question, we need to set the values of $\tilde{\gamma}_t$. A reasonable choices to fix $\tilde{\gamma}_t = \kappa \tau_t$, for some constant κ , since τ_t can be thought as the “effective noise level” at iteration t . There is a one-to-one correspondence between κ and the regularization parameter λ in the LASSO [17]. We thus define the function

$$G(\tau^2; \sigma^2) \equiv \sigma^2 + \frac{1}{\delta} \mathbb{E} \left\{ (\eta(\Theta + \tau Z; \kappa \tau) - \Theta)^2 \right\}, \quad (129)$$

which of course depends implicitly on p_Θ , κ , δ . State evolution is then the one-dimensional recursion $\tau_{t+1}^2 = G(\tau_t^2; \sigma^2)$. For the sequence τ_t to stay bounded we assume $\lim_{\tau^2 \rightarrow \infty} G(\tau^2; \sigma^2)/\tau^2 < 1$ which can always be ensured by taking κ sufficiently large.

Let us first consider the noiseless case $\sigma = 0$. Since $G(0; 0) = 0$, we know that $\tau = 0$ is always a fixed point. It is not hard to shown [16] that indeed $\lim_{t \rightarrow \infty} \tau_t^2 = 0$ if and

only if this is the unique non-negative fixed point, see figure below. If this condition is satisfied, AMP reconstructs exactly the signal θ , and due to the correspondence with the LASSO, also basis pursuit (the LASSO with $\lambda \rightarrow 0$) reconstructs exactly θ .



Notice that this condition is sharp: If it is not satisfied, then AMP and the LASSO fail to reconstruct θ , despite vanishing noise. In order to derive the phase transition location, remember that by the definition of minimax risk of soft thresholding, cf. Section 3.1, we have, assuming $\kappa = \ell(\varepsilon)$ to be set in the optimal way

$$G(\tau^2; 0) = \frac{1}{\delta} \mathbb{E} \{ (\eta(\Theta + Z; \kappa\tau) - \Theta)^2 \} \leq \frac{M(\varepsilon)}{\delta} \tau^2. \quad (130)$$

Hence $\tau_{t+1}^2 \leq (M(\varepsilon)/\delta)\tau_t^2$ and, if

$$\delta > M(\varepsilon), \quad (131)$$

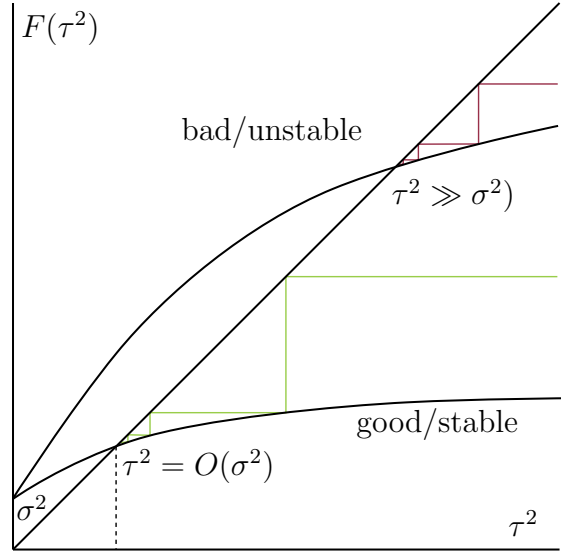
then $\tau_t^2 \rightarrow 0$ and AMP (LASSO) reconstructs θ with vanishing error. This bound is in fact tight: For $\delta < M(\varepsilon)$, any probability distribution p_Θ with $p_\Theta(\{0\}) = 1 - \varepsilon$, and any threshold parameter κ , the mean square error remains bounded away from zero.

Recalling the definition of $\delta = n/p$, the condition $\delta > M(\varepsilon)$ corresponds to requiring a sufficient number of samples, as compared to the sparsity. It is interesting to recover the very sparse regime from this point of view. Recall from previous lectures that $M(\varepsilon) \approx 2\varepsilon \log(1/\varepsilon)$ for small ε . The condition $\delta > M(\varepsilon)$ then translates to $\delta \gtrsim 2\varepsilon \log(1/\varepsilon)$ or, in other words, $(n/p) \gtrsim 2(s_0/p) \log(p/s_0)$. Thus, we obtain the condition –already discussed before– that the number of samples must be as large as the number of non-zero coefficients, times a logarithmic factor. Reconstruction is possible if and only if

$$n \gtrsim 2s_0 \log \frac{p}{s_0}, \quad (132)$$

a condition that we have seen in previous lectures.

In the noisy case, we cannot hope to achieve perfect reconstruction. In this case, we say that estimation is stable if there is a constant C such that, for any $\theta \in \mathbb{R}^p$, $R(\theta; \hat{\theta}) \leq C\sigma^2$. This setting is sketched in the figure below. Exact reconstruction at $\sigma = 0$ translate into a fixed point $\tau_*^2 = O(\sigma^2)$ and hence stability. Inexact reconstruction corresponds to a fixed point of order 1 and hence lack of stability.



Again by choosing a suitable threshold value κ , we can ensure that the minimax bound (130) is valid and hence

$$\tau_{t+1}^2 \leq \sigma^2 + \frac{M(\varepsilon)}{\delta} \tau_t^2. \quad (133)$$

Taking the limit $t \rightarrow \infty$, in the case that $\delta > M(\varepsilon)$ we have that

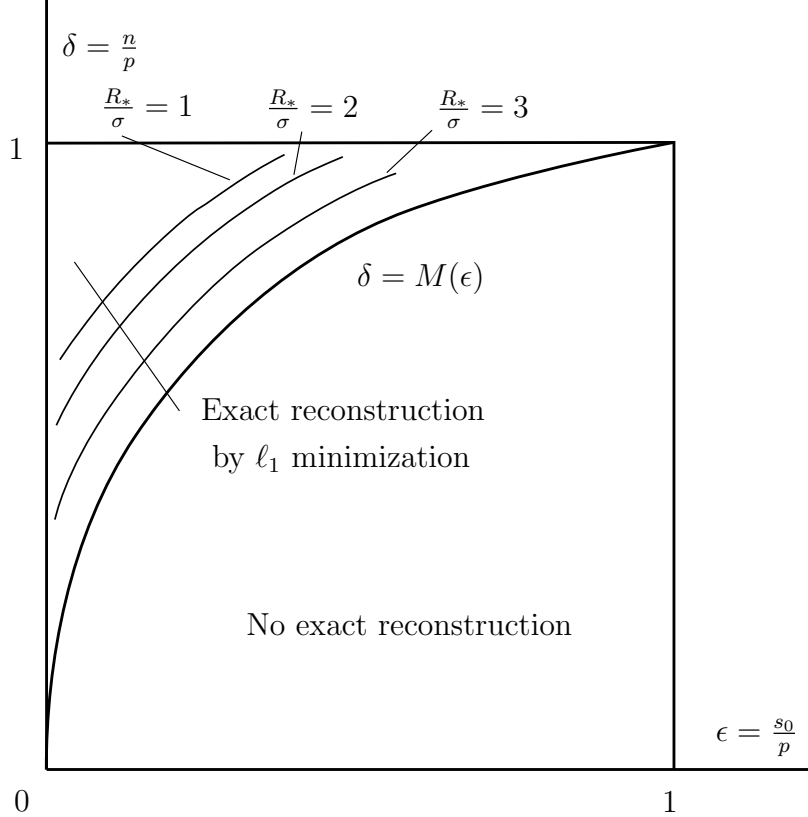
$$\tau_*^2 \leq \frac{\sigma^2}{1 - (M(\varepsilon)/\delta)}. \quad (134)$$

This establishes that the following is an upper bound on the asymptotic mean square error of AMP, and hence (by the equivalence discussed above) of the LASSO,

$$R_\infty(\theta; \hat{\theta}) = \begin{cases} \frac{M(\varepsilon)}{\delta - M(\varepsilon)} \sigma^2, & \text{if } M(\varepsilon) < \delta, \\ \infty & \text{otherwise.} \end{cases} \quad (135)$$

As proven in [17, 42], this result holds indeed with equality (these papers have slightly different normalizations of the noise variance σ^2).

A qualitative sketch of resulting phase diagram in ε and δ is in the figure below. As anticipated above, if $\delta > M(\varepsilon)$, i.e. in the regime in which exact reconstruction is feasible through basis pursuit in zero noise, reconstruction is also stable with respect to noise.



Again, let us consider the sparse regime $\epsilon \rightarrow 0$. Assuming $M(\epsilon) \ll \delta$, and substituting $M(\epsilon) \approx 2\epsilon \log(1/\epsilon)$ together with the definitions of ϵ and δ we get

$$R_\infty(\theta; \hat{\theta}) = \sigma^2 \frac{M(\epsilon)}{\delta - M(\epsilon)} \approx \frac{\sigma^2}{\delta} 2\epsilon \log \frac{1}{\epsilon} = \frac{s_0 \sigma^2}{n} 2 \log \frac{p}{s_0}. \quad (136)$$

We therefore rederived the same behavior already established in the previous section under the RIP assumption. Apart from the factor $2 \log(p/s_0)$ the risk is the same ‘as if’ we knew the support of θ .

6. THE HIDDEN CLIQUE PROBLEM

One of the most surprising facts about sparse regression is that we can achieve ideal estimation error, using a low complexity algorithm, namely by solving a convex optimization problem such as the LASSO. Indeed –at first sight– one might have suspected it necessary to search over possible supports of size s_0 , a task that requires at least $\binom{p}{s_0}$ operations, and is therefore non-polynomial. Unfortunately, this is not always the case. There are problems in which a huge gap exists between the statistical limits of estimation (i.e. the minimax risk achieved by an arbitrary estimator) and the computational limits (i.e. the minimax risk achieved by any estimator computable in polynomial time). The hidden clique (or hidden submatrix) problem is a prototypical example of this class of computationally hard estimation problems. Recently, reductions to this problem were used to prove that other estimation problems are hard as well [56].

We next define the problem. Let Q_0 and Q_1 be two given probability distributions on \mathbb{R} . For a set $S \subseteq \{1, 2, \dots, n\}$ we let $\mathbf{W} \in \mathbb{R}^{n \times n}$ be a symmetric random matrix with entries $(\mathbf{W}_{ij})_{i \leq j}$ independent, with distribution:

$$\mathbf{W}_{ij} \sim Q_1, \quad \text{if } i, j \in S \tag{137}$$

$$\mathbf{W}_{ij} \sim Q_0, \quad \text{otherwise.} \tag{138}$$

The problem is to find the set S given one realization of \mathbf{W} .

Example 1: Suppose $Q_0 = \mathcal{N}(0, 1)$ and $Q_1 = \mathcal{N}(\mu, 1)$ to be two Gaussian distributions with different means and same known variance (which we set –without loss of generality– equal to one). The model is then equivalent to the following

$$\mathbf{W} = \mu u_S u_S^\top + \mathbf{Z}, \tag{139}$$

where u_S is the indicator vector of the set S , namely $(u_S)_i = 1$ if $i \in S$, and $(u_S)_i = 0$ otherwise.

Example 2: This is the original setting of the hidden clique problem from [57]. Both Q_0 and Q_1 are Bernoulli distributions:

$$Q_0 = \frac{1}{2} \delta_{-1} + \frac{1}{2} \delta_{+1}, \tag{140}$$

$$Q_1 = \delta_{+1}. \tag{141}$$

There is a straightforward way to interpret this as a graph problem. Let G be the random graph on n vertices $\{1, \dots, n\}$ whereby two vertices i, j are joined by an edge if and only if $\mathbf{W}_{ij} = +1$. Then G is an Erdős-Renyi random graph (with edge density $1/2$) to which a clique has been added with support on S .


For simplicity of exposition, we will focus for the rest of this lecture on the Bernoulli case, i.e. on the last example above. We will use interchangeably the language of random graphs and the one of random matrices. All of our results can in fact be generalized to arbitrary probability distributions Q_0, Q_1 under suitable tail conditions, as shown in [58].

We will denote by $k = |S|$ the size of the hidden set. It is not hard to see that the problem is easy for k large (both from the statistical and the computational point of view), and hard for k small (both computationally and a statistically). Indeed, for k sufficiently large, a simple degree based heuristics is successful. This is based on the remark that vertices in the clique have a slightly higher degree than others. Hence sorting the vertices by degree, the first k vertices should provide a good estimate of S .


Proposition 1. *Let \hat{S} be the set of k vertices with larges degree in G . If $k \geq \sqrt{(2 + \varepsilon)n \log n}$, then with high probability $\hat{S} = S$.*

Proof. Let D_i denote the degree of vertex i . If $i \notin S$, then $D_i \sim \text{Binom}(n - 1, 1/2)$. In particular, standard concentration bounds on independent random variables yield $\mathbb{P}\{D_i \geq \mathbb{E}D_i + t\} \leq \exp(-2t^2/n)$. By a union bound (the same already used to analyze denoising in Section 3.1), and using $\mathbb{E}(D_i) = (n - 1)/2$, we have, for any $\varepsilon' > 0$, with probability converging to one as $n \rightarrow \infty$,


$$\max_{i \notin S} D_i \leq \frac{n}{2} + \sqrt{(1 + \varepsilon') \frac{n \log n}{2}}. \tag{142}$$



(a) A random graph with a planted clique.



(b) The same graph, but with the vertices shuffled.



(c) Retrieving the clique in the shuffled graph

On the other hand, if $i \in S$, then $D_i \sim k - 1 + \text{Binom}(n - k, 1/2)$. Hence, by a similar union bound

$$\min_{i \in S} D_i \geq \frac{n + k}{2} - \sqrt{(1 + \varepsilon') \frac{n \log k}{2}}. \quad (143)$$

The claim follows by using together the above, and selecting a suitable value ε' . \square

For k too small, the problem becomes statistically intractable because the planted clique is not the unique clique of size k . Hence no estimator can distinguish between the set S and another set S' that supports a different (purely random) clique. The next theorem characterizes this statistical threshold.

Proposition 2. *Let $\varepsilon > 0$ be fixed. Then, for $k < 2(1 - \varepsilon) \log_2 n$ any estimator \widehat{S} is such that $\widehat{S} \neq S$ with probability converging to one as $n \rightarrow \infty$.*

Viceversa, for $k < 2(1 - \varepsilon) \log_2 n$ there exists an estimator \widehat{S} such that $\widehat{S} = S$ with probability converging to one as $n \rightarrow \infty$.

Proof. We will not present a complete proof but only sketch the fundamental reason for a threshold $k \approx 2 \log_2 n$ and leave to the reader the task of filling the details.

The basic observation is that the largest ‘purely random’ clique is of size approximately $2 \log_2 n$. As a consequence, for k larger than this threshold, searching for a clique of size k returns the planted clique.

More precisely, let $\mathcal{G}(n, 1/2)$ be an Erdős-Renyi random graph with edge density $1/2$ (i.e. a random graph where each edge is present independently with probability $1/2$). We will show that the largest clique in $\mathcal{G}(n, 1/2)$ is with high probability of size between $2(1 - \varepsilon) \log_2 n$ and $2(1 + \varepsilon) \log_2 n$.

This claim can be proved by a moment calculation. In particular, for proving that the largest clique cannot be much larger than $2 \log_2 n$, it is sufficient to compute the expected number of cliques of size ℓ . Letting $N(\ell; n)$ denote the number of cliques of size ℓ in $\mathcal{G}(n, 1/2)$, we have

$$\mathbb{E} N(\ell; n) = \binom{n}{\ell} 2^{-\binom{\ell}{2}} \approx n^\ell 2^{-\ell^2/2} = 2^{\ell \log_2 n - \ell^2/2}. \tag{144}$$

For $\ell > 2(1 + \varepsilon) \log_2 n$ the exponent is negative and the expectation vanishes as $n \rightarrow \infty$. In fact $\sum_{\ell \geq 2(1+\varepsilon) \log_2 n} \mathbb{E} N(\ell; n)$ vanishes as well. By Markov inequality, it follows that –with high probability– no clique has size larger than $2(1 + \varepsilon) \log_2 n$. \square

The catch with the last proposition is that the estimator needs not to be computable in polynomial time. Indeed the estimator implicitly assumed in the proof requires searching over all subsets of k vertices, which takes time at least $\binom{n}{k} \approx n^k$. For k above the threshold, this is $\exp\{c(\log n)^2\}$, that is super-polynomial.

To summarize, with unlimited computational resources we can find planted cliques as soon as their size is larger than $c \log_2 n$ for any $c > 2$. This is the fundamental statistical barrier towards estimating the set S . On the other hand, the naive degree-based heuristic described above, correctly identifies the clique if $k \geq \sqrt{cn \log n}$. There is a huge gap between the fundamental statistical limit, and what is achieved by a simple polynomial-time algorithm. This begs the question as to whether this gap can be filled by more advanced algorithmic ideas.

A key observation, due to Alon, Krivelevich and Sudakov [59] is that the matrix \mathbf{W} –in expectation– a rank-one matrix. Namely

$$\mathbb{E}\{\mathbf{W}\} = u_S u_S^{sT}, \tag{145}$$

and therefore S can be reconstructed from the eigenvalue decomposition of $\mathbb{E}\{\mathbf{W}\}$. Of course $\mathbb{E}\{\mathbf{W}\}$ is not available, but one can hope the random part of \mathbf{W} not to perturb too much the leading eigenvector. In other words, one can compute the principal eigenvector $v_1(\mathbf{W})$, i.e. the eigenvector of \mathbf{W} with largest eigenvalue, and use its largest entries to estimate the clique. For instance, one can take the k vertices corresponding to the entries of $v_1(\mathbf{W})$ with largest absolute value.

This spectral approach allows to reduce the minimum detectable clique size by a factor $\sqrt{\log n}$, with respect to the degree heuristics of 1.

Theorem 3 (Alon, Krivelevich and Sudakov, 1998). *There exists an algorithm that returns an estimate \hat{S} of the set S , with the same complexity as computing the principal eigenvector of \mathbf{W} , and such that the following holds. If $k > 100\sqrt{n}$, then $\hat{S} = S$ with probability converging to one as $n \rightarrow \infty$.*

Proof sketch. Again, we will limit ourselves to explaining the basic argument. The actual proof requires some additional steps.

Then the matrix \mathbf{W} has the form

$$\mathbf{W} = u_S u_S^\top + \mathbf{Z} - \mathbf{Z}_{S,S}, \quad (146)$$

where \mathbf{Z} is a Wigner matrix i.e. a matrix with i.i.d. zero-mean entries $(\mathbf{Z}_{ij})_{i \leq j}$, and $\mathbf{Z}_{S,S}$ is the restriction of \mathbf{Z} to indices in S . In the present case, the entries distribution is Bernoulli

$$\mathbf{Z}_{ij} = \begin{cases} +1, & \text{with probability } 1/2 \\ -1, & \text{with probability } 1/2 \end{cases}. \quad (147)$$

By the celebrated Füredi-Komlos theorem of [60], the operator norm of this matrix (i.e. the maximum of the largest eigenvalue of \mathbf{Z} and the largest eigenvalue of $-\mathbf{Z}$) is upper bounded as $\|\mathbf{Z}\|_2 \leq (2 + \varepsilon)\sqrt{n}$, with high probability. By the same argument $\|\mathbf{Z}_{S,S}\|_2 \leq (2 + \varepsilon)\sqrt{k}$, which is much smaller than $\|\mathbf{Z}\|_2$.

We view \mathbf{W} as a perturbation of the matrix $u_S u_S^\top$ (whose principal, normalized, eigenvector is u_S/\sqrt{k}). Matrix perturbation theory implies that the largest eigenvector is perturbed by an amount proportional to the norm of the perturbation and inversely proportional to the gap between top eigenvalue and second eigenvalue of the perturbed matrix. More precisely, Davis-Kahan ‘sin theta’ theorem yields (for $v_1 = v_1(\mathbf{W})$)

$$\sin \theta(v_1, u_S) \leq \frac{\|\mathbf{Z} - \mathbf{Z}_{S,S}\|_2}{\lambda_1(u_S u_S^\top) - \lambda_2(\mathbf{W})}, \quad (148)$$

where $\lambda_\ell(\mathbf{A})$ denotes the ℓ -th largest eigenvalue of matrix \mathbf{A} , and $\theta(a, b)$ is the angle between vectors a and b . We of course have $\lambda_1(u_S u_S^\top) = \|u_S\|_2^2 = k$, and $\lambda_2(\mathbf{W}) \leq \lambda_2(u_S u_S^\top) + \|\mathbf{Z} - \mathbf{Z}_{S,S}\|_2$. Therefore, for

$$\sin \theta(v_1, u_S) \leq \frac{\|\mathbf{Z} - \mathbf{Z}_{S,S}\|_2}{k - \|\mathbf{Z} - \mathbf{Z}_{S,S}\|_2} \quad (149)$$

$$\leq \frac{2.1\sqrt{n}}{k - 2.1\sqrt{n}} \leq \text{frac145.}, \quad (150)$$

where the last inequality holds with high probability by Füredi-Komlos theorem. Using standard trigonometry, this bound can be immediately converted in a bound on the ℓ_2 distance between $v_1(\mathbf{W})$ and the unperturbed eigenvector:

$$\left\| v_1 - \frac{u_S}{\sqrt{k}} \right\|_2 \leq \frac{1}{40}. \quad (151)$$

We can then select the set B of k vertices that correspond to the k entries of v_1 with largest absolute value. The last bound does not guarantee that B coincide with S , but it implies that B must have a substantial overlap with S . The estimator \widehat{S} is constructed by selecting the k vertices in $\{1, 2, \dots, p\}$ that have the largest number of neighbors in B . \square

It is useful to pause for a few remarks on this result.

Remark 1. *The complexity of the above algorithm is the same as the one of computing the principal eigenvector $v_1(\mathbf{W})$. Under the assumptions of the theorem, this is non-degenerate and in fact, there is a large gap between the first eigenvalue and the second one, say $\max(\lambda_2(\mathbf{W}), |\lambda_n(\mathbf{W})|) \leq (1/2)\lambda_1(\mathbf{W})$.*

Hence, $v_1(\mathbf{W})$ can be computed efficiently through power iteration, i.e. by computing the sequence of vectors $v^{(t+1)} = \mathbf{W}v^{(t)}$. Each operation takes at most n^2 operations, and due to the fast convergence, $O(\log n)$ iterations are sufficient for implementing the above algorithm. We will revisit power iteration in the following.

Remark 2. *The eigenvalues and eigenvectors of a random matrix of the form (146) have been studied in detail in statistics (under the name of ‘spiked model’) and probability theory (as ‘low-rank perturbation of Wigner matrices’), see e.g. [61, 62, 63]. These work unveil a phase transition phenomenon that, in the present application, can be stated as follows. Assume $k, n \rightarrow \infty$ with $k/\sqrt{n} = \kappa \in (0, \infty)$. Then*

$$\lim_{n \rightarrow \infty} |\langle v_1(\mathbf{W}), u_S/\sqrt{k} \rangle| = \begin{cases} 0 & \text{if } \kappa \leq 1, \\ \sqrt{1 - \kappa^{-2}} & \text{otherwise.} \end{cases} \quad (152)$$

In other words, for $k \leq (1 - \varepsilon)\sqrt{n}$ the principal eigenvector of \mathbf{W} is essentially uncorrelated with the hidden set S . The barrier at k of order \sqrt{n} is not a proof artifact, but instead a fundamental limit related to this phase transition.

On the other hand, a more careful analysis of the spectral method can possibly show that it succeeds for all $k \geq (1 + \varepsilon)\sqrt{n}$. (Here and above $\varepsilon > 0$ is an arbitrary constant).

Remark 3. *A clever trick by Alon and collaborators [59], allow to find cliques of size $k \geq \delta\sqrt{n}$ for any fixed constant $\delta > 0$ in polynomial time. The price to pay is that the computational complexity increases rapidly as δ gets smaller. More precisely, we can identify sets of size $k \geq \delta\sqrt{n}$ for any with time complexity of order $n^{O(\log(1/\delta))}$.*

To see this, we use the spectral method as a routine that is able to find the clique with high probability provided $k \geq c\sqrt{n}$ for some constant c . First assume that an oracle gives us one node in the clique. We can solve the problem with $k \gtrsim c\sqrt{n}/2$. Indeed we can focus our attention on the set of neighbors of the node provided by the oracle. There is about $n/2$ such neighbors, and they contain a clique of size $k - 1$, hence the spectral method will succeed under the stated condition.

We then observe that we do not need an such: we can search for the vertex that the oracle would tell us by blowing up the runtime by a factor at most n (indeed only a \sqrt{n} factor is sufficient, since one every \sqrt{n} vertices is in the clique). In this way we can trade a factor of $\sqrt{2}$ in k by an n -fold increase of the runtime. This construction can be repeated $O(\log(1/\delta))$ times to achieve the trade-off mentioned above.

6.1. An iterative thresholding approach. Throughout this section, we shall normalize the data and work with the matrix $\mathbf{A} = \mathbf{W}/\sqrt{n}$. As we saw in the previous section, the principal eigenvector of \mathbf{A} carries important information about the set S , and in particular it is correlated with the indicator vector u_S , if the hidden set S is large enough. Also an efficient way to compute the principal eigenvector is through power iteration

$$v^{(t+1)} = \mathbf{A} v^{(t)}. \quad (153)$$

Note that the resulting vector $v^{(t)}$ will not –in general– be sparse, if not, a posteriori, because of the correlation with u_S . It is therefore a natural idea to modify the power iteration by introducing a non-linearity that enforces sparsity:

$$\theta^{(t+1)} = \mathbf{A} f_t(\theta^{(t)}), \quad (154)$$

where $\theta \in \mathbb{R}^n$ and $f_t : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a non-linear function that enforces sparsity. To be definite, we will assume throughout that the initialization is $\theta^{(0)} = (1, 1, \dots, 1)$, the all-ones vector.

For ease of exposition, we shall focus on separable functions and denote by f_t the action of this function on each component. In other words, with a slight abuse of notation, we will write $f_t(v) = (f_t(v_1), f_t(v_2), \dots, f_t(v_n))$ when $v = (v_1, v_2, \dots, v_n)$. Example of such a function might be

- *Positive soft thresholding:* $f_t(x) = (x - \lambda_t)_+$ for some iteration-dependent threshold λ_t . The threshold can be chosen so that, on average $f_t(\theta^{(t)})$ has a number of non-zeros of order k .
- *Positive hard thresholding:* $f_t(x) = x \mathbb{1}(x \geq \lambda_t)$ (here $\mathbb{1}$ is the indicator function: $\mathbb{1}(B) = 1$ if B is true and $= 0$ otherwise). Again λ_t is a threshold.
- *Logistic nonlinearity:*

$$f_t(x) = \frac{1}{1 + \exp(-a_t(x - \lambda_t))}, \quad (155)$$

where λ_t plays the role of a ‘soft threshold.’

Which function should we choose? Which thresholds? Will this approach beat the simple power iteration (i.e. $f_t(x) = x$)?

In order to address these questions, we will carry out a simple heuristic analysis of the above non-linear power iteration. Remarkably, we will see in the next section that this analysis yields the correct answer for a modified version of the same algorithm –a message passing algorithms. Our discussion is based on [58], and we refer to that paper for all omitted details, formal statements and derivations.

The heuristic analysis requires to consider separately vertices in S and outside S :

- (1) For $i \notin S$, the non-linear power iteration (154) reads

$$\theta_i^{(t+1)} = \sum_{j=1}^n \mathbf{A}_{ij} f_t(\theta_j^{(t)}). \quad (156)$$

Since in this case the variables $\{\mathbf{A}_{ij}\}_{j \in [n]}$ are i.i.d. with mean zero and variance $1/n$, it is natural to guess –by central limit theorem– $\theta_i^{(t+1)}$ to be approximately normal with mean 0 and variance $(1/n) \sum_{j=1}^n f_t^2(\theta_j^{(t)})$. Repeating this argument inductively, we conclude that that $\theta_i^t \sim \mathcal{N}(0, \sigma_t^2)$, where –by the law of large numbers applied to $(1/n) \sum_{j=1}^n f_t^2(\theta_j^{(t)})$ –

$$\sigma_{t+1}^2 = \mathbb{E}\{f_t(\sigma_t Z)^2\}, \quad (157)$$

where the expectation is taken with respect to $Z \sim \mathcal{N}(0, 1)$. The initialization $\theta^{(0)} = u$ implies $\sigma_1^2 = f_0(1)^2$

- (2) For $i \in S$, we have $\mathbf{A}_{ij} = \kappa$ if $j \in S$ as well, and $\mathbf{A}_{ij} = \mathbf{Z}_{ij}/\sqrt{n}$ having zero mean and variance $1/n$ otherwise. Hence

$$\theta_i^{(t+1)} = \kappa \sum_{j \in S} f_t(\theta_j^{(t)}) + \frac{1}{\sqrt{n}} \sum_{j \in [n] \setminus S} \mathbf{Z}_{ij} f_t(\theta_j^{(t)}). \quad (158)$$

By the same argument as above, the second part gives rise to a zero-mean Gaussian contribution, with variance σ_t^2 , and the first has non-zero mean and negligible variance. We conclude that $\theta_i^{(t)}$ is approximately $\mathcal{N}(\mu_t, \sigma_t^2)$ with σ_t given recursively by Eq. (157). Applying the law of large numbers to the non-zero mean contribution, we get the recursion

$$\mu_{t+1} = \kappa \mathbb{E}\{f_t(\mu_t + \sigma_t Z)\}, \quad (159)$$

where the expectation is taken with respect to $Z \sim \mathcal{N}(0, 1)$, and the initialization $\theta^{(0)} = u$ implies $\mu_1 = \kappa f_0(1)$ (recall that κ is defined as the limit of k/\sqrt{n}).

A few important remarks.

The above derivation is of course incorrect! The problem is that the central limit theorem cannot be applied to the right-hand side of Eq. (156) because the summands are not independent. Indeed, each term $f_t(\theta_j^{(t)})$ depends on all the entries of the matrix \mathbf{A} .

The conclusion that we reached is incorrect. It is not true that, asymptotically, $\theta_i^{(t)}$ is approximately Gaussian, with the above mean and variance.

Surprisingly, the conclusion is correct for a slightly modified algorithm, namely a message passing algorithm that will be introduced in the next section. This is a highly non-trivial phenomenon

6.2. A message passing algorithm. We modify the non-linear power iteration (156) by transforming it into a message passing algorithm, whose underlying graph is the complete graph with n vertices. The iteration variables are ‘messages’ $\theta_{i \rightarrow j}^{(t)}$ for each $i \neq j$ (with $\theta_{i \rightarrow j}^{(t)} \neq \theta_{j \rightarrow i}^{(t)}$). These are updated using the rule

$$\theta_{i \rightarrow j}^{(t)} = \sum_{k \in [n] \setminus j} \mathbf{A}_{ik} f_t(\theta_{k \rightarrow i}^{(t)}). \quad (160)$$

The only difference with respect to the iteration (156) is that we exclude the term $k = j$ from the sum. Despite this seemingly negligible change (one out of n terms is dropped), the statistical properties of this algorithm are significantly different from the ones of the nonlinear power iteration (156), even in the limit $n \rightarrow \infty$. In particular, the Gaussian limit derived heuristically in the previous section, holds for the message passing algorithm. Informally, we have, as $n \rightarrow \infty$,

$$\theta_{i \rightarrow j}^{(t)} \sim \begin{cases} \mathcal{N}(\mu_t, \sigma_t^2) & \text{if } i \in S, \\ \mathcal{N}(0, \sigma_t^2) & \text{if } i \notin S, \end{cases} \quad (161)$$

where μ_t, σ_t are determined by the state evolution equations (157) and (157).

Let us stress that we did not yet choose the functions $f_t(\cdot)$: we defer this choice, as well as an analysis of state evolution to the next section. Before this, we note that –as in the case of sparse regression– an approximate message passing (AMP) version

of this algorithm can be derived by writing $\theta_{i \rightarrow j}^{(t)} = \theta_i^{(t)} + \delta \theta_{i \rightarrow j}^{(t)}$ and linearizing in the latter correction. This calculation leads to the simple AMP iteration

$$\theta^{(t+1)} = \mathbf{A} f_t(\theta^{(t)}) - \mathbf{b}_t f_{t-1}(\theta^{(t-1)}), \quad (162)$$

where the ‘Onsager term’ \mathbf{b}_t is given in this case by

$$\mathbf{b}_t = \frac{1}{n} \sum_{i=1}^n f'_t(\theta_i^{(t)}). \quad (163)$$

6.3. Analysis and optimal choice of $f_t(\cdot)$. We now consider the implications of state evolution for the performance of the above message passing algorithms. For the sake of simplicity, we will refer to the AMP form (162), but analogous statements hold for the message passing version (160). Informally, state evolution implies that

$$\theta^{(t)} \approx \mu_t u_S + \sigma_t z, \quad (164)$$

where $z \sim \mathcal{N}(0, \mathbf{I}_n)$, and this statement holds asymptotically in the sense of finite-dimensional marginals.

In other words, we can interpret $\theta^{(t)}$ as a noisy observation of the unknown vector u_S , corrupted by Gaussian noise. This suggest to choose $f_t(\cdot)$ as the posterior expectation denoiser. Namely, for $y \in \mathbb{R}$


$$f_t^{\text{opt}}(y) = \mathbb{E}\{U \mid \mu_t U + \sigma_t Z = y\}, \quad (165)$$

where $U \sim \text{Bernoulli}(p)$ for $p = k/n = \kappa/\sqrt{n}$, and $Z \sim \mathcal{N}(0, 1)$ independently of U . A simple calculation yields the explicit expression

$$f_t^{\text{opt}}(y) = \frac{\delta}{\delta + (1 - \delta) \exp\left(-\frac{\mu_t}{\sigma_t^2} y + \frac{\mu_t^2}{\sigma_t^2}\right)}. \quad (166)$$

This is indeed empirically the best choice for the non-linearity $f_t^{\text{opt}}(\cdot)$. We shall next rederive it from a different point of view, which also allow to characterize its behavior.

Reconsider again the Gaussian limit (164). It is clear that the quality of the information contained in $\theta^{(t)}$ depends on signal to noise ratio μ_t/σ_t . Note that u_S is very sparse, hence the vector $\theta^{(t)}$ is undistinguishable from a zero-mean Gaussian vector unless $(\mu_t/\sigma_t) \rightarrow \infty$. Indeed, unless this happens, the entries $\theta_i^{(t)}$, $i \in S$, are hidden in the tail of the zero-mean entries $\theta_i^{(t)}$, $i \in [n] \setminus S$, see figure below. It turns out that, by optimally choosing $f_t(\cdot)$, this happens if and only if $\kappa > 1/\sqrt{e}$. In other words, the message passing algorithm succeeds with high probability as long as k is larger than $(1 + \varepsilon)\sqrt{n/e}$, for any $\varepsilon > 0$.



In order to determine whether $\tilde{\mu}_t \equiv (\mu_t/\sigma_t) \rightarrow \infty$, note that –without loss of generality– we can rescale the functions $f_t(\cdot)$ so that $\sigma_t = 1$ for all t (simply replacing $f_t(z)$ by $f_t(z)/\mathbb{E}\{f_t(Z)^2\}^{1/2}$ in Eq. (160), or in Eq. (162)). After this normalization, Eq. (159) yields

$$\tilde{\mu}_{t+1} = \kappa \frac{\mathbb{E}\{f_t(\tilde{\mu}_t + Z)\}}{\mathbb{E}\{f_t(Z)^2\}^{1/2}}. \quad (167)$$

Note that

$$\mathbb{E}\{f_t(\tilde{\mu}_t + Z)\} = \int f_t(x) \frac{1}{\sqrt{2\pi}} e^{-(x-\mu_t)^2/2} dx \quad (168)$$

$$= e^{-\tilde{\mu}_t^2/2} \mathbb{E}\{f_t(Z) e^{\tilde{\mu}_t x}\} \leq e^{\tilde{\mu}_t^2} \mathbb{E}\{f_t(Z)^2\}^{1/2}. \quad (169)$$

where the last inequality follows from Cauchy-Schwartz inequality. The inequality is saturated by taking $f_t(x) = e^{\tilde{\mu}_t x - (\tilde{\mu}_t^2/2)}$, that yields the state evolution recursion

$$\tilde{\mu}_{t+1} = \kappa e^{\tilde{\mu}_t^2/2}. \quad (170)$$

It is immediate to study this recursion, and conclude that $\tilde{\mu}_t \rightarrow \infty$ if and only if $\kappa > 1/\sqrt{e}$.

The above analysis indeed yields the correct threshold for a message passing algorithm, as proved in [58]. (For proving the theorem below, a ‘cleaning’ step is added to the message passing algorithm.)

Theorem 4 (Deshpande, Montanari, 2014). *There exists an algorithm with time complexity $O(n^2 \log n)$, that outputs an estimate \hat{S} such that –if $k > (1 + \epsilon)\sqrt{n/e}$ – then $\hat{S} = S$ with probability converging to one as $n \rightarrow \infty$.*

In other words, the message passing algorithm is able to find cliques smaller by a factor $1/\sqrt{e}$ with respect to spectral methods, with no increase in complexity. A natural research question is the following:

Is it possible to planted find cliques of size $(1-\epsilon)\sqrt{n/e}$ in time $O(n^2 \log n)$?

The paper [58] provides a partially positive answer to this question, by showing that no ‘local algorithm’ (a special class of linear-time algorithm) can beat message passing algorithms for a sparse-graph version of the planted clique problem.

Let us conclude by showing how the last derivation agrees in fact with the guess (166) for the optimal non-linearity. Note that $\delta = \kappa/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$. In this limit

$$f_t^{\text{opt}}(y) \approx C_t \exp \left\{ \frac{\mu_t}{\sigma_t^2} y \right\}. \quad (171)$$

This coincides with the choice optimizing the state evolution threshold, once we set $\sigma_t = 1$ (that entails no loss of generality).

REFERENCES

- [1] L. Wasserman, *All of statistics: a concise course in statistical inference*. Springer, 2004.
- [2] ———, *All of nonparametric statistics*. Springer, 2006.
- [3] I. M. Johnstone, *Gaussian estimation: Sequence and wavelet models*, 2011, draft version, December 27, 2011, available at <http://statweb.stanford.edu/~imj/GE12-27-11.pdf>.
- [4] A. B. Tsybakov and V. Zaiats, *Introduction to nonparametric estimation*. Springer, 2009, vol. 11.
- [5] P. Bühlmann and S. Van De Geer, *Statistics for high-dimensional data: methods, theory and applications*. Springer, 2011.
- [6] F. J. Herrmann, M. P. Friedlander, and Ö. Yilmaz, “Fighting the curse of dimensionality: Compressive sensing in exploration seismology,” *Signal Processing Magazine*, vol. 29, no. 3, pp. 88–100, 2012.
- [7] D. Donoho, I. Johnstone, and A. Montanari, “Accurate prediction of phase transitions in compressed sensing via a connection to minimax denoising,” *Information Theory, IEEE Transactions on*, vol. 59, no. 6, pp. 3396–3433, June 2013.
- [8] C. F. Gauss, *Theoria combinationis observationum erroribus minimis obnoxiae*. H. Dieterich, 1823.
- [9] W. James and C. Stein, “Estimation with quadratic loss,” in *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, vol. 1, no. 1961, 1961, pp. 361–379.
- [10] D. L. Donoho and I. M. Johnstone, “Minimax risk over l_p balls,” *Prob. Th. and Rel. Fields*, vol. 99, pp. 277–303, 1994.
- [11] ———, “Neo-classical minimax problems, thresholding, and adaptation,” *Bernoulli*, pp. 39–62, 1994.
- [12] D. Donoho and I. Johnstone, “Adapting to unknown smoothness via wavelet shrinkage,” *J. Amer. Statist. Assn.*, vol. 90, pp. 1200–1224, 1995.
- [13] D. L. Donoho and I. M. Johnstone, “Minimax estimation via wavelet shrinkage,” *Annals of Statistics*, vol. 26, pp. 879–921, 1998.
- [14] B. Derrida, “Random-energy model: An exactly solvable model of disordered systems,” *Physical Review B*, vol. 24, no. 5, p. 2613, 1981.
- [15] D. L. Donoho and J. M. Johnstone, “Ideal spatial adaptation by wavelet shrinkage,” *Biometrika*, vol. 81, no. 3, pp. 425–455, 1994.
- [16] D. L. Donoho, A. Maleki, and A. Montanari, “Message-passing algorithms for compressed sensing,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 45, pp. 18 914–18 919, 2009.
- [17] ———, “The noise-sensitivity phase transition in compressed sensing,” *Information Theory, IEEE Transactions on*, vol. 57, no. 10, pp. 6920–6941, 2011.
- [18] R. Tibshirani, “Regression shrinkage and selection with the Lasso,” *J. Royal. Statist. Soc B*, vol. 58, pp. 267–288, 1996.
- [19] S. Chen and D. Donoho, “Examples of basis pursuit,” in *Proceedings of Wavelet Applications in Signal and Image Processing III*, San Diego, CA, 1995.

- [20] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic decomposition by basis pursuit,” *SIAM journal on scientific computing*, vol. 20, no. 1, pp. 33–61, 1998.
- [21] A. Juditsky and A. Nemirovski, “First order methods for nonsmooth convex large-scale optimization,” *Optimization for Machine Learning*, pp. 121–148, 2011.
- [22] I. Daubechies, M. Defrise, and C. De Mol, “An iterative thresholding algorithm for linear inverse problems with a sparsity constraint,” *Communications on pure and applied mathematics*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [23] A. Beck and M. Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems,” *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.
- [24] E. J. Candés and T. Tao, “Decoding by linear programming,” *IEEE Trans. on Inform. Theory*, vol. 51, pp. 4203–4215, 2005.
- [25] E. Candés and T. Tao, “The dantzig selector: Statistical estimation when p is much larger than n ,” *The Annals of Statistics*, pp. 2313–2351, 2007.
- [26] P. J. Bickel, Y. Ritov, and A. B. Tsybakov, “Simultaneous analysis of Lasso and Dantzig selector,” *Amer. J. of Mathematics*, vol. 37, pp. 1705–1732, 2009.
- [27] M. Rudelson and R. Vershynin, “On sparse reconstruction from fourier and gaussian measurements,” *Communications on Pure and Applied Mathematics*, vol. 61, no. 8, pp. 1025–1045, 2008.
- [28] J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, and R. G. Baraniuk, “Beyond nyquist: Efficient sampling of sparse bandlimited signals,” *Information Theory, IEEE Transactions on*, vol. 56, no. 1, pp. 520–544, 2010.
- [29] D. L. Donoho, “High-dimensional centrally symmetric polytopes with neighborliness proportional to dimension,” *Discrete & Computational Geometry*, vol. 35, no. 4, pp. 617–652, 2006.
- [30] D. L. Donoho and J. Tanner, “Neighborliness of randomly projected simplices in high dimensions,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, no. 27, pp. 9452–9457, 2005.
- [31] —, “Sparse nonnegative solution of underdetermined linear equations by linear programming,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, no. 27, pp. 9446–9451, 2005.
- [32] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, “The convex geometry of linear inverse problems,” *Foundations of Computational Mathematics*, vol. 12, no. 6, pp. 805–849, 2012.
- [33] E. Candés and B. Recht, “Simple bounds for recovering low-complexity models,” *Mathematical Programming*, vol. 141, no. 1-2, pp. 577–589, 2013.
- [34] M. Stojnic, “A framework to characterize performance of lasso algorithms,” *arXiv preprint arXiv:1303.7291*, 2013.
- [35] S. Oymak, C. Thrampoulidis, and B. Hassibi, “The squared-error of generalized lasso: A precise analysis,” *arXiv preprint arXiv:1311.0830*, 2013.
- [36] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp, “Living on the edge: A geometric theory of phase transitions in convex optimization,” *arXiv preprint arXiv:1303.6672*, 2013.
- [37] Y. Kabashima, T. Wadayama, and T. Tanaka, “A typical reconstruction limit for compressed sensing based on l_p -norm minimization,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2009, no. 09, p. L09003, 2009.
- [38] S. Rangan, V. Goyal, and A. K. Fletcher, “Asymptotic analysis of map estimation via the replica method and compressed sensing,” in *Advances in Neural Information Processing Systems*, 2009, pp. 1545–1553.
- [39] D. Guo, D. Baron, and S. Shamai, “A single-letter characterization of optimal noisy compressed sensing,” in *Communication, Control, and Computing, 2009. Allerton 2009. 47th Annual Allerton Conference on*. IEEE, 2009, pp. 52–59.
- [40] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová, “Statistical-physics-based reconstruction in compressed sensing,” *Physical Review X*, vol. 2, no. 2, p. 021005, 2012.
- [41] M. Bayati and A. Montanari, “The dynamics of message passing on dense graphs, with applications to compressed sensing,” *Information Theory, IEEE Transactions on*, vol. 57, no. 2, pp. 764–785, 2011.

- [42] —, “The lasso risk for gaussian matrices,” *Information Theory, IEEE Transactions on*, vol. 58, no. 4, pp. 1997–2017, 2012.
- [43] A. Javanmard and A. Montanari, “State evolution for general approximate message passing algorithms, with applications to spatial coupling,” *Information and Inference*, p. iat004, 2013.
- [44] D. Donoho, A. Javanmard, and A. Montanari, “Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing,” *Information Theory, IEEE Transactions on*, vol. 59, no. 11, pp. 7434–7464, Nov 2013.
- [45] M. Bayati, M. Lelarge, and A. Montanari, “Universality in polytope phase transitions and message passing algorithms,” *Annals of Applied Probability*, 2014.
- [46] D. L. Donoho, A. Maleki, and A. Montanari, “Message passing algorithms for compressed sensing: I. motivation and construction,” in *Information Theory Workshop (ITW), 2010 IEEE*. IEEE, 2010, pp. 1–5.
- [47] T. Richardson and R. Urbanke, *Modern coding theory*. Cambridge University Press, 2008.
- [48] M. Mezard and A. Montanari, *Information, physics, and computation*. Oxford University Press, 2009.
- [49] S. Som and P. Schniter, “Compressive imaging using approximate message passing and a markov-tree prior,” *Signal Processing, IEEE Transactions on*, vol. 60, no. 7, pp. 3439–3448, 2012.
- [50] S. Rangan, “Generalized Approximate Message Passing for Estimation with Random Linear Mixing,” in *IEEE Intl. Symp. on Inform. Theory*, St. Petersburg, Aug. 2011, pp. 2168 – 2172.
- [51] D. Donoho and A. Montanari, “High dimensional robust m-estimation: Asymptotic variance via approximate message passing,” *arXiv preprint arXiv:1310.7320*, 2013.
- [52] C. A. Metzler, A. Maleki, and R. G. Baraniuk, “From denoising to compressed sensing,” *arXiv preprint arXiv:1406.4175*, 2014.
- [53] J. Tan, Y. Ma, and D. Baron, “Compressive imaging via approximate message passing with image denoising,” *arXiv preprint arXiv:1405.4429*, 2014.
- [54] J. Barbier and F. Krzakala, “Replica analysis and approximate message passing decoder for superposition codes,” *arXiv preprint arXiv:1403.8024*, 2014.
- [55] E. Bolthausen, “An iterative construction of solutions of the tap equations for the sherrington–kirkpatrick model,” *Communications in Mathematical Physics*, vol. 325, no. 1, pp. 333–366, 2014.
- [56] Q. Berthet and P. Rigollet, “Complexity theoretic lower bounds for sparse principal component detection,” in *Conference on Learning Theory*, 2013, pp. 1046–1066.
- [57] M. Jerrum, “Large cliques elude the metropolis process,” *Random Structures & Algorithms*, vol. 3, no. 4, pp. 347–359, 1992.
- [58] Y. Deshpande and A. Montanari, “Finding Hidden Cliques of Size \sqrt{rtN}/e in Nearly Linear Time,” *arXiv:1304.7047*, 2013.
- [59] N. Alon, M. Krivelevich, and B. Sudakov, “Finding a large hidden clique in a random graph,” in *Proceedings of the ninth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 1998, pp. 594–598.
- [60] Z. Füredi and J. Komlós, “The eigenvalues of random symmetric matrices,” *Combinatorica*, vol. 1, no. 3, pp. 233–241, 1981.
- [61] D. Féral and S. Péché, “The largest eigenvalue of rank one deformation of large wigner matrices,” *Communications in mathematical physics*, vol. 272, no. 1, pp. 185–228, 2007.
- [62] M. Capitaine, C. Donati-Martin, and D. Féral, “The largest eigenvalues of finite rank deformation of large wigner matrices: convergence and nonuniversality of the fluctuations,” *The Annals of Probability*, vol. 37, no. 1, pp. 1–47, 2009.
- [63] —, “Central limit theorems for eigenvalues of deformations of wigner matrices,” in *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, vol. 48, no. 1. Institut Henri Poincaré, 2012, pp. 107–133.