
Contents

List of Figures	xi
Preface	xiii
Acknowledgements	xvii
Notational Conventions	xix
About Proofs	xxi
On Computational Tractability	xxi
1 Sparse Recovery via ℓ_1 Minimization	1
1.1 Compressed Sensing: What is it about?	1
1.1.1 Signal Recovery Problem	1
1.1.2 Signal Recovery: Parametric and nonparametric cases	2
1.1.3 Compressed Sensing via ℓ_1 minimization: Motivation	6
1.2 Validity of sparse signal recovery via ℓ_1 minimization	8
1.2.1 Validity of ℓ_1 minimization in the noiseless case	8
1.2.2 Imperfect ℓ_1 minimization	11
1.2.3 Regular ℓ_1 recovery	13
1.2.4 Penalized ℓ_1 recovery	14
1.2.5 Discussion	14
1.3 Verifiability and tractability issues	19
1.3.1 Restricted Isometry Property and s -goodness of random matrices	20
1.3.2 Verifiable sufficient conditions for $\mathbf{Q}_q(s, \kappa)$	20
1.3.3 Tractability of $\mathbf{Q}_\infty(s, \kappa)$	22
1.4 Exercises for Chapter 1	26
1.5 Proofs	30
1.5.1 Proofs of Theorem 1.3, 1.4	30
1.5.2 Proof of Theorem 1.5	32
1.5.3 Proof of Proposition 1.7	33
1.5.4 Proof of Propositions 1.8 and 1.12	36
1.5.5 Proof of Proposition 1.10	37
1.5.6 Proof of Proposition 1.13	39
2 Hypothesis Testing	41
2.1 Preliminaries from Statistics: Hypotheses, Tests, Risks	41
2.1.1 Hypothesis Testing Problem	41
2.1.2 Tests	42
2.1.3 Testing from repeated observations	42
2.1.4 Risk of a simple test	45

2.1.5	Two-point lower risk bound	46
2.2	Hypothesis Testing via Euclidean Separation	49
2.2.1	Situation	49
2.2.2	Pairwise Hypothesis Testing via Euclidean Separation	50
2.2.3	Euclidean Separation, Repeated Observations, and Majority Tests	55
2.2.4	From Pairwise to Multiple Hypotheses Testing	58
2.3	Detectors and Detector-Based Tests	65
2.3.1	Detectors and their risks	65
2.3.2	Detector-based tests	65
2.4	Simple observation schemes	72
2.4.1	Simple observation schemes—Motivation	72
2.4.2	Simple observation schemes—The definition	73
2.4.3	Simple observation schemes—Examples	74
2.4.4	Simple observation schemes—Main result	79
2.4.5	Simple observation schemes—Examples of optimal detectors	83
2.5	Testing multiple hypotheses	87
2.5.1	Testing unions	88
2.5.2	Testing multiple hypotheses “up to closeness”	91
2.5.3	Illustration: Selecting the best among a family of estimates	97
2.6	Sequential Hypothesis Testing	105
2.6.1	Motivation: Election polls	105
2.6.2	Sequential hypothesis testing	108
2.6.3	Concluding remarks	113
2.7	Measurement Design in simple observation schemes	113
2.7.1	Motivation: Opinion polls revisited	113
2.7.2	Measurement Design: Setup	115
2.7.3	Formulating the MD problem	116
2.8	Affine detectors beyond simple observation schemes	123
2.8.1	Situation	124
2.8.2	Main result	132
2.9	Beyond the scope of affine detectors: lifting the observations	139
2.9.1	Motivation	139
2.9.2	Quadratic lifting: Gaussian case	140
2.9.3	Quadratic lifting—Does it help?	142
2.9.4	Quadratic lifting: Sub-Gaussian case	145
2.9.5	Generic application: Quadratically constrained hypotheses	147
2.10	Exercises for Chapter 2	157
2.10.1	Two-point lower risk bound	157
2.10.2	Around Euclidean Separation	157
2.10.3	Hypothesis testing via ℓ_1 -separation	157
2.10.4	Miscellaneous exercises	163
2.11	Proofs	168
2.11.1	Proof of the observation in Remark 2.8	168
2.11.2	Proof of Proposition 2.6 in the case of quasi-stationary K -repeated observations	168
2.11.3	Proof of Theorem 2.23	172
2.11.4	Proof of Proposition 2.37	175
2.11.5	Proof of Proposition 2.43	176
2.11.6	Proof of Proposition 2.46	180

3	From Hypothesis Testing to Estimating Functionals	185
3.1	Estimating linear forms on unions of convex sets	185
3.1.1	The problem	186
3.1.2	The estimate	187
3.1.3	Main result	189
3.1.4	Near-optimality	190
3.1.5	Illustration	191
3.2	Estimating N -convex functions on unions of convex sets	193
3.2.1	Outline	194
3.2.2	Estimating N -convex functions: Problem setting	197
3.2.3	Bisection estimate: Construction	199
3.2.4	Building Bisection estimate	201
3.2.5	Bisection estimate: Main result	202
3.2.6	Illustration	203
3.2.7	Estimating N -convex functions: An alternative	205
3.3	Estimating linear forms beyond simple observation schemes	211
3.3.1	Situation and goal	212
3.3.2	Construction and main results	213
3.3.3	Estimation from repeated observations	216
3.3.4	Application: Estimating linear forms of sub-Gaussianity parameters	218
3.4	Estimating quadratic forms via quadratic lifting	222
3.4.1	Estimating quadratic forms, Gaussian case	222
3.4.2	Estimating quadratic form, sub-Gaussian case	228
3.5	Exercises for Chapter 3	238
3.6	Proofs	250
3.6.1	Proof of Proposition 3.3	250
3.6.2	Verifying 1-convexity of the conditional quantile	253
3.6.3	Proof of Proposition 3.4	254
3.6.4	Proof of Proposition 3.14	258
4	Signal Recovery by Linear Estimation	260
	Overview	260
4.1	Preliminaries: Executive summary on Conic Programming	262
4.1.1	Cones	262
4.1.2	Conic problems and their duals	263
4.1.3	Schur Complement Lemma	265
4.2	Near-optimal linear estimation from Gaussian observations	265
4.2.1	Situation and goal	265
4.2.2	Building a linear estimate	267
4.2.3	Byproduct on semidefinite relaxation	274
4.3	From ellitopes to spectratopes	275
4.3.1	Spectratopes: Definition and examples	275
4.3.2	Semidefinite relaxation on spectratopes	277
4.3.3	Linear estimates beyond ellitopic signal sets and $\ \cdot\ _2$ -risk	278
4.4	Linear estimates of stochastic signals	291
4.4.1	Minimizing Euclidean risk	293
4.4.2	Minimizing $\ \cdot\ $ -risk	294
4.5	Linear estimation under uncertain-but-bounded noise	295
4.5.1	Uncertain-but-bounded noise	295

4.5.2	Mixed noise	299
4.6	Calculus of ellitopes/spectratopes	300
4.7	Exercises for Chapter 4	302
4.7.1	Linear estimates vs. Maximum Likelihood	302
4.7.2	Measurement Design in Signal Recovery	303
4.7.3	Around semidefinite relaxation	306
4.7.4	Around Propositions 4.4 and 4.14	317
4.7.5	Signal recovery in Discrete and Poisson observation schemes	335
4.7.6	Numerical lower-bounding minimax risk	347
4.7.7	Around \mathcal{S} -Lemma	359
4.7.8	Miscellaneous exercises	360
4.8	Proofs	361
4.8.1	Preliminaries	361
4.8.2	Proof of Proposition 4.6	364
4.8.3	Proof of Proposition 4.8	366
4.8.4	Proof of Lemma 4.17	368
4.8.5	Proofs of Propositions 4.5, 4.16 and 4.19	371
4.8.6	Proofs of Propositions 4.18 and 4.19, and justification of Remark 4.20	383
5	Signal Recovery Beyond Linear Estimates	386
	Overview	386
5.1	Polyhedral estimation	386
5.1.1	Motivation	386
5.1.2	Generic polyhedral estimate	388
5.1.3	Specifying sets \mathcal{H}_δ for basic observation schemes	390
5.1.4	Efficient upper-bounding of $\mathfrak{R}[H]$ and contrast design, I. . .	392
5.1.5	Efficient upper-bounding of $\mathfrak{R}[H]$ and contrast design, II. .	399
5.1.6	Assembling estimates: Contrast aggregation	411
5.1.7	Numerical illustration	413
5.1.8	Calculus of compatibility	413
5.2	Recovering signals from nonlinear observations by Stochastic Optimization	415
5.2.1	Problem setting	415
5.2.2	Assumptions	417
5.2.3	Estimating via Sample Average Approximation	420
5.2.4	Stochastic Approximation estimate	423
5.2.5	Numerical illustration	425
5.2.6	“Single-observation” case	428
5.3	Exercises for Chapter 5	431
5.3.1	Estimation by Stochastic Optimization	431
5.4	Proofs	440
5.4.1	Proof of (5.8)	440
5.4.2	Proof of Lemma 5.6	441
5.4.3	Verification of (5.44)	442
5.4.4	Proof of Proposition 5.10	443
	Solutions to Selected Exercises	447
6.1	Solutions for Chapter 1	447
6.2	Solutions for Chapter 2	454

6.2.1	Two-point lower risk bound	454
6.2.2	Around Euclidean Separation	455
6.2.3	Hypothesis testing via ℓ_1 separation	457
6.2.4	Miscellaneous exercises	465
6.3	Solutions for Chapter 3	477
6.4	Solutions for Chapter 4	495
6.4.1	Linear Estimates vs. Maximum Likelihood	495
6.4.2	Measurement Design in Signal Recovery	497
6.4.3	Around semidefinite relaxation	502
6.4.4	Around Propositions 4.4 and 4.14	518
6.4.5	Numerical lower-bounding minimax risk	572
6.4.6	Around \mathcal{S} -Lemma	586
6.4.7	Miscellaneous exercises	589
6.5	Solutions for Chapter 5	592
6.5.1	Estimation by Stochastic Optimization	592
Appendix: Executive Summary on Efficient Solvability of Convex Optimization Problems		609
Bibliography		613
Index		629

Chapter One

Sparse Recovery via ℓ_1 Minimization

In this chapter, we overview basic results of *Compressed Sensing*, a relatively new and rapidly developing area in Statistics and Signal Processing dealing with recovering signals (vectors x from some \mathbf{R}^n) from their noisy observations $Ax + \eta$ (A is a given $m \times n$ *sensing matrix*, η is observation noise) in the case when the number of observations m is much smaller than the signal's dimension n , but is essentially larger than the “true” dimension—the number of nonzero entries—in the signal. This setup leads to a deep, elegant and highly innovative theory and possesses quite significant application potential. It should be added that along with the plain sparsity (small number of nonzero entries), Compressed Sensing deals with other types of “low-dimensional structure” hidden in high-dimensional signals, most notably, with the case of *low rank matrix recovery*—when the signal is a matrix, and sparse signals are matrices with low ranks—and the case of *block sparsity*, where the signal is a block vector, and sparsity means that only a small number of blocks are nonzero. In our presentation, we do *not* consider these extensions, and restrict ourselves to the simplest sparsity paradigm.

1.1 COMPRESSED SENSING: WHAT IS IT ABOUT?

1.1.1 Signal Recovery Problem

One of the basic problems in Signal Processing is the problem of recovering a *signal* $x \in \mathbf{R}^n$ from noisy observations

$$y = Ax + \eta \tag{1.1}$$

of a linear image of the signal under a given *sensing mapping* $x \mapsto Ax : \mathbf{R}^n \rightarrow \mathbf{R}^m$; in (1.1), η is the *observation error*. Matrix A in (1.1) is called *sensing matrix*.

Recovery problems of the outlined types arise in many applications, including, but *by far* not reducing to,

- *communications*, where x is the signal sent by the transmitter, y is the signal recorded by the receiver, and A represents the communication channel (reflecting, e.g., dependencies of decays in the signals' amplitude on the transmitter-receiver distances); η here typically is modeled as the standard (zero mean, unit covariance matrix) m -dimensional Gaussian noise;¹

¹While the “physical” noise indeed is often Gaussian with zero mean, its covariance matrix is not necessarily the unit matrix. Note, however, that a zero mean Gaussian noise η always can be represented as $Q\xi$ with standard Gaussian ξ . Assuming that Q is known and is nonsingular (which indeed is so when the covariance matrix of η is positive definite), we can rewrite (1.1) equivalently as

$$Q^{-1}y = [Q^{-1}A]x + \xi$$

and treat $Q^{-1}y$ and $Q^{-1}A$ as our new observation and new sensing matrix; the new observation

- *image reconstruction*, where the signal x is an image—a 2D array in the usual photography, or a 3D array in tomography—and y is data acquired by the imaging device. Here η in many cases (although not always) can again be modeled as the standard Gaussian noise;
- *linear regression*, arising in a wide range of applications. In linear regression, one is given m pairs “input $a^i \in \mathbf{R}^n$ ” to a “black box,” with output $y_i \in \mathbf{R}$. Sometimes we have reason to believe that the output is corrupted by noise version of the “existing in nature,” but unobservable, “ideal output” $y_i^* = x^T a^i$ which is just a linear function of the input (this is called “linear regression model,” with inputs a^i called “regressors”). Our goal is to convert actual observations (a^i, y_i) , $1 \leq i \leq m$, into estimates of the *unknown* “true” vector of parameters x . Denoting by A the matrix with the rows $[a^i]^T$ and assembling individual observations y_i into a single observation $y = [y_1; \dots; y_m] \in \mathbf{R}^m$, we arrive at the problem of recovering vector x from noisy observations of Ax . Here again the most popular model for η is the standard Gaussian noise.

1.1.2 Signal Recovery: Parametric and nonparametric cases

Recovering signal x from observation y would be easy if there were no observation noise ($\eta = 0$) and the rank of matrix A were equal to the dimension n of the signals. In this case, which arises only when $m \geq n$ (“more observations than unknown parameters”), and is typical in this range of m and n , the desired x would be the unique solution to the system of linear equations, and to find x would be a simple problem of Linear Algebra. Aside from this trivial “enough observations, no noise” case, people over the years have looked at the following two versions of the recovery problem:

Parametric case: $m \gg n$, η is nontrivial noise with zero mean, say, standard Gaussian. This is the classical statistical setup with the emphasis on how to use numerous available observations in order to suppress in the recovery, to the extent possible, the influence of observation noise.

Nonparametric case: $m \ll n$.² If addressed literally, this case seems to be senseless: when the number of observations is less than the number of unknown parameters, even in the noiseless case we arrive at the necessity to solve an undetermined (fewer equations than unknowns) system of linear equations. Linear Algebra says that if solvable, the system has infinitely many solutions. Moreover, the solution set (an affine subspace of positive dimension) is unbounded, meaning that the solutions are in no sense close to each other. A typical way to make the case of $m \ll n$ meaningful is to add to the observations (1.1) some a priori information about the signal. In traditional Nonparametric Statistics, this additional information is summarized in a *bounded convex set* $X \subset \mathbf{R}^n$, given to us in advance, known to contain the true signal x . This set usually is such that *every signal* $x \in X$ can be approximated by a linear combination of $s = 1, 2, \dots, n$ vectors

noise ξ is indeed standard. Thus, in the case of Gaussian zero mean observation noise, to assume the noise standard Gaussian is the same as to assume that its covariance matrix is known.

²Of course, this is a blatant simplification—the nonparametric case covers also a variety of important and by far nontrivial situations in which m is comparable to n or larger than n (or even $\gg n$). However, this simplification is very convenient, and we will use it in this introduction.

from a properly selected basis known to us in advance (“dictionary” in the slang of signal processing) within accuracy $\delta(s)$, where $\delta(s)$ is a function, known in advance, approaching 0 as $s \rightarrow \infty$. In this situation, with appropriate A (e.g., just the unit matrix, as in the denoising problem), we can select some $s \ll m$ and try to recover x as if it were a vector from the linear span E_s of the first s vectors of the outlined basis [54, 86, 124, 112, 208]. In the “ideal case,” $x \in E_s$, recovering x in fact reduces to the case where the dimension of the signal is $s \ll m$ rather than $n \gg m$, and we arrive at the well-studied situation of recovering a signal of low (compared to the number of observations) dimension. In the “realistic case” of x $\delta(s)$ -close to E_s , deviation of x from E_s results in an additional component in the recovery error (“bias”); a typical result of traditional Nonparametric Statistics quantifies the resulting error and minimizes it in s [86, 124, 178, 222, 223, 230, 239]. Of course, this outline of the traditional approach to “nonparametric” (with $n \gg m$) recovery problems is extremely sketchy, but it captures the most important fact in our context: with the traditional approach to nonparametric signal recovery, one assumes that after representing the signals by vectors of their coefficients in properly selected base, the n -dimensional signal to be recovered can be well approximated by an s -sparse (at most s nonzero entries) signal, with $s \ll n$, and this sparse approximation can be obtained by zeroing out all but the first s entries in the signal vector. The assumption just formulated indeed takes place for signals obtained by discretization of smooth uni- and multivariate functions, and this class of signals for several decades was the main, if not the only, focus of Nonparametric Statistics.

Compressed Sensing. The situation changed dramatically around the year 2000 as a consequence of important theoretical breakthroughs due to D. Donoho, T. Tao, J. Romberg, E. Candes, and J.-J. Fuchs, among many other researchers [49, 44, 45, 46, 48, 67, 68, 69, 70, 93, 94]; as a result of these breakthroughs, a novel and rich area of research, called *Compressed Sensing*, emerged.

In the Compressed Sensing (CS) setup of the Signal Recovery problem, as in the traditional Nonparametric Statistics approach to the $m \ll n$ case, it is assumed that after passing to an appropriate basis, the signal to be recovered is s -sparse (has $\leq s$ nonzero entries, with $s \ll m$), or is well approximated by an s -sparse signal. The difference with the traditional approach is that now we assume *nothing* about the location of the nonzero entries. Thus, the a priori information about the signal x both in the traditional and in the CS settings is summarized in a set X known to contain the signal x we want to recover. The difference is that in the traditional setting, X is a bounded convex and “nice” (well approximated by its low-dimensional cross-sections) set, while in CS this set is, computationally speaking, a “monster”: already in the simplest case of recovering *exactly* s -sparse signals, X is the union of all s -dimensional coordinate planes, which is a heavily combinatorial entity.

Note that, in many applications we indeed can assume that the true vector of parameters x is sparse. Consider, e.g., the following story about signal detection. *There are n locations where signal transmitters could be placed, and m locations with the receivers. The contribution of a signal of unit magnitude originating in location j to the signal measured by receiver i is a known quantity A_{ij} , and signals originating in different locations merely sum up in the receivers. Thus, if x is the n -dimensional vector with entries x_j representing the magnitudes of signals transmitted in locations $j = 1, 2, \dots, n$, then the m -dimensional vector y of measurements of the m receivers is $y =$*

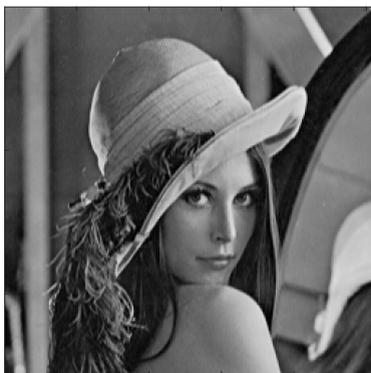
$Ax + \eta$, where η is the observation noise. Given y , we intend to recover x .

Now, if the receivers are, say, hydrophones registering noises emitted by submarines in a certain part of the Atlantic, tentative positions of “submarines” being discretized with resolution 500 m, the dimension of the vector x (the number of points in the discretization grid) may be in the range of tens of thousands, if not tens of millions. At the same time, presumably, there is only a handful of “submarines” (i.e., nonzero entries in x) in the area.

To “see” sparsity in everyday life, look at the 256×256 image at the top of Figure 1.1. The image can be thought of as a $256^2 = 65,536$ -dimensional vector comprised of the pixels’ intensities in gray scale, and there is not much sparsity in this vector. However, when representing the image in the *wavelet basis*, whatever it means, we get a “nearly sparse” vector of wavelet coefficients (this is true for typical “non-pathological” images). At the bottom of Figure 1.1 we see what happens when we zero out all but a small percentage of the wavelet coefficients largest in magnitude and replace the true image by its sparse—in the wavelet basis—approximations.

This simple visual illustration along with numerous similar examples shows the “everyday presence” of sparsity and the possibility to utilize it when compressing signals. The difficulty, however, is that simple compression—compute the coefficients of the signal in an appropriate basis and then keep, say, 10% of the largest in magnitude coefficients—requires us to start with digitalizing the signal—representing it as an array of all its coefficients in some orthonormal basis. These coefficients are inner products of the signal with vectors of the basis; for a “physical” signal, like speech or image, these inner products are computed by analog devices, with subsequent discretization of the results. After the measurements are discretized, processing the signal (denoising, compression, storing, etc.) can be fully computerized. The major (to some extent, already actualized) advantage of Compressed Sensing is in the possibility to reduce the “analogous effort” in the outlined process: instead of computing analogously n linear forms of n -dimensional signal x (its coefficients in a basis), we use an analog device to compute $m \ll n$ other linear forms of the signal and then use the signal’s sparsity in a basis known to us in order to recover the signal reasonably well from these m observations.

In our “picture illustration” this technology would work (in fact, works—it is called “single pixel camera” [83]; see Figure 1.2) as follows: in reality, the digital 256×256 image on the top of Figure 1.1 was obtained by an analog device—a digital camera which gets on input an analog signal (light of varying intensity along the field of view caught by camera’s lens) and discretizes the light’s intensity in every pixel to get the digitalized image. We then can compute the wavelet coefficients of the digitalized image, compress its representation by keeping, say, just 10% of leading coefficients, etc., but “the damage is already done”—we have already spent our analog resources to get the entire digitalized image. The technology utilizing Compressed Sensing would work as follows: instead of measuring and discretizing light intensity in each of the 65,536 pixels, we compute (using an analog device) the integral, taken over the field of view, of the product of light intensity and an analog-generated “mask.” We repeat it for, say, 20,000 different masks, thus obtaining measurements of 20,000 linear forms of our 65,536-dimensional signal. Next we utilize, via the Compressed Sensing machinery, the signal’s sparsity in the wavelet basis in order to recover the signal from these 20,000 measurements. With this approach, we reduce the “analog component” of signal processing effort,



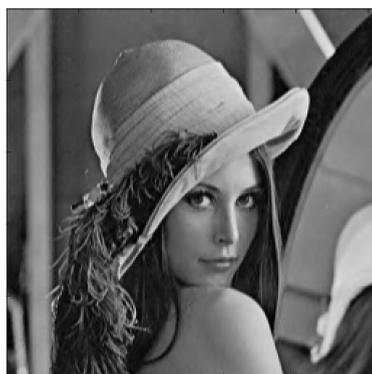
1% of leading wavelet coefficients (97.83 % of energy) kept



5% of leading wavelet coefficients (99.51 % of energy) kept



10% of leading wavelet coefficients (99.82% of energy) kept



25% of leading wavelet coefficients (99.97% of energy) kept

Figure 1.1: Top: true 256×256 image; bottom: sparse in the wavelet basis approximations of the image. Wavelet basis is orthonormal, and a natural way to quantify near-sparsity of a signal is to look at the fraction of total energy (sum of squares of wavelet coefficients) stored in the leading coefficients; these are the “energy data” presented in the figure.

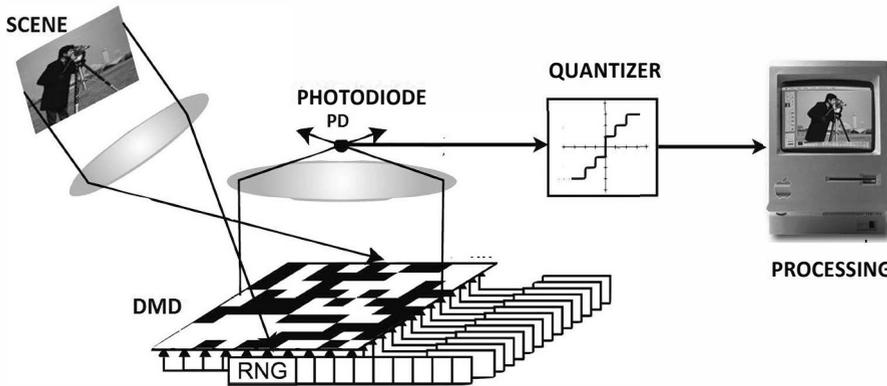


Figure 1.2: Single-pixel camera.

at the price of increasing the “computerized component” of the effort (instead of ready-to-use digitalized image directly given by 65,536 analog measurements, we need to recover the image by applying computationally nontrivial decoding algorithms to our 20,000 “indirect” measurements). When taking pictures with your camera or iPad, the game is not worth the candle—the analog component of taking usual pictures is cheap enough, and decreasing it at the cost of nontrivial decoding of the digitalized measurements would be counterproductive. There are, however, important applications where the advantages stemming from reduced “analog effort” outweigh significantly the drawbacks caused by the necessity to use nontrivial computerized decoding [96, 176].

1.1.3 Compressed Sensing via ℓ_1 minimization: Motivation

1.1.3.1 Preliminaries

In principle there is nothing surprising in the fact that under reasonable assumption on the $m \times n$ sensing matrix A we may hope to recover from noisy observations of Ax an s -sparse signal x , with $s \ll m$. Indeed, assume for the sake of simplicity that there are no observation errors, and let $\text{Col}_j[A]$ be j -th column in A . If we knew the locations $j_1 < j_2 < \dots < j_s$ of the nonzero entries in x , identifying x could be reduced to solving the system of linear equations $\sum_{\ell=1}^s x_{i_\ell} \text{Col}_{j_\ell}[A] = y$ with m equations and $s \ll m$ unknowns; assuming every s columns in A to be linearly independent (a quite unrestrictive assumption on a matrix with $m \geq s$ rows), the solution to the above system is unique, and is exactly the signal we are looking for. Of course, the assumption that we know the locations of nonzeros in x makes the recovery problem completely trivial. However, it suggests the following course of action: given noiseless observation $y = Ax$ of an s -sparse signal x , let us solve the

combinatorial optimization problem

$$\min_z \{ \|z\|_0 : Az = y \}, \quad (1.2)$$

where $\|z\|_0$ is the number of nonzero entries in z . Clearly, the problem has a solution with the value of the objective at most s . Moreover, it is immediately seen that if every $2s$ columns in A are linearly independent (which again is a very unrestrictive assumption on the matrix A provided that $m \geq 2s$), then the true signal x is the unique optimal solution to (1.2).

What was said so far can be extended to the case of noisy observations and “nearly s -sparse” signals x . For example, assuming that the observation error is “uncertain-but-bounded,” specifically some known norm $\|\cdot\|$ of this error does not exceed a given $\epsilon > 0$, and that the true signal is s -sparse, we could solve the combinatorial optimization problem

$$\min_z \{ \|z\|_0 : \|Az - y\| \leq \epsilon \}. \quad (1.3)$$

Assuming that every $m \times 2s$ submatrix \bar{A} of A is not just with linearly independent columns (i.e., with trivial kernel), but is reasonably well conditioned,

$$\|\bar{A}w\| \geq C^{-1}\|w\|_2$$

for all $(2s)$ -dimensional vectors w , with some constant C , it is immediately seen that the true signal x underlying the observation and the optimal solution \hat{x} of (1.3) are close to each other within accuracy of order of ϵ : $\|x - \hat{x}\|_2 \leq 2C\epsilon$. It is easily seen that the resulting error bound is basically as good as it could be.

We see that the difficulties with recovering sparse signals stem not from the lack of information; they are of purely computational nature: (1.2) is a difficult combinatorial problem. As far as known theoretical complexity guarantees are concerned, they are not better than “brute force” search through all guesses on where the nonzeros in x are located—by inspecting first the only option that there are no nonzeros in x at all, then by inspecting n options that there is only one nonzero, for every one of n locations of this nonzero, then $n(n-1)/2$ options that there are exactly two nonzeros, etc., until the current option results in a solvable system of linear equations $Az = y$ in variables z with entries restricted to vanish outside the locations prescribed by the current option. The running time of this “brute force” search, beyond the range of small values of s and n (by far too small to be of any applied interest), is by many orders of magnitude larger than what we can afford in reality.³

A partial remedy is as follows. Well, if we do not know how to minimize the “bad” objective $\|z\|_0$ under linear constraints, as in (1.2), let us “approximate” this objective with one which we do know how to minimize. The true objective is separable: $\|z\|_0 = \sum_{i=1}^n \xi(z_i)$, where $\xi(s)$ is the function on the axis equal to 0 at the origin and equal to 1 otherwise. As a matter of fact, the separable functions which

³When $s = 5$ and $n = 100$, a sharp upper bound on the number of linear systems we should process before termination in the “brute force” algorithm is $\approx 7.53e7$ —a lot, but perhaps doable. When $n = 200$ and $s = 20$, the number of systems to be processed jumps to $\approx 1.61e27$, which is by many orders of magnitude beyond our “computational grasp”; we would be unable to carry out that many computations even if the fate of the mankind were at stake. And from the perspective of Compressed Sensing, $n = 200$ still is a completely toy size, 3–4 orders of magnitude less than we would like to handle.

we do know how to minimize under linear constraints are sums of *convex* functions of z_1, \dots, z_n . The most natural candidate to the role of *convex* approximation of $\xi(s)$ is $|s|$; with this approximation, (1.2) converts into the ℓ_1 *minimization problem*

$$\min_z \left\{ \|z\|_1 := \sum_{i=1}^n |z_i| : Az = y \right\}, \quad (1.4)$$

and (1.3) becomes the convex optimization problem

$$\min_z \{ \|z\|_1 : \|Az - y\| \leq \epsilon \}. \quad (1.5)$$

Both problems are efficiently solvable, which is nice; the question, however, is how relevant these problems are in our context—whether it is true that they do recover the “true” s -sparse signals in the noiseless case, or “nearly recover” these signals when the observation error is small. Since we want to be able to handle *any* s -sparse signal, the validity of ℓ_1 recovery—its ability to recover well *every* s -sparse signal—depends solely on the sensing matrix A . Our current goal is to understand which sensing matrices are “good” in this respect.

1.2 VALIDITY OF SPARSE SIGNAL RECOVERY VIA ℓ_1 MINIMIZATION

What follows is based on the standard basic results of Compressed Sensing theory originating from [19, 49, 45, 44, 46, 47, 48, 67, 69, 70, 93, 94, 232] and augmented by the results of [129, 130, 132, 133].⁴

1.2.1 Validity of ℓ_1 minimization in the noiseless case

The minimal requirement on sensing matrix A which makes ℓ_1 minimization valid is to guarantee the correct recovery of *exactly* s -sparse signals in the *noiseless* case, and we start with investigating this property.

1.2.1.1 Notational convention

From now on, for a vector $x \in \mathbf{R}^n$

- $I_x = \{j : x_j \neq 0\}$ stands for the *support* of x ; we also set

$$I_x^+ = \{j : x_j > 0\}, I_x^- = \{j : x_j < 0\} \quad [\Rightarrow I_x = I_x^+ \cup I_x^-];$$

- for a subset I of the index set $\{1, \dots, n\}$, x_I stands for the vector obtained from x by zeroing out entries with indices *not* in I , and I^c for the complement of I :

$$I^c = \{i \in \{1, \dots, n\} : i \notin I\};$$

- for $s \leq n$, x^s stands for the vector obtained from x by zeroing out all but the s

⁴In fact, in the latter source, an extension of the sparsity, the so-called block sparsity, is considered; in what follows, we restrict the results of [130] to the case of plain sparsity.

- entries largest in magnitude.⁵ Note that x^s is the best s -sparse approximation of x in all ℓ_p norms, $1 \leq p \leq \infty$;
- for $s \leq n$ and $p \in [1, \infty]$, we set

$$\|x\|_{s,p} = \|x^s\|_p;$$

note that $\|\cdot\|_{s,p}$ is a norm.

1.2.1.2 s -Goodness

Definition of s -goodness. Let us say that an $m \times n$ sensing matrix A is s -good if whenever the true signal x underlying *noiseless* observations is s -sparse, this signal will be recovered *exactly* by ℓ_1 minimization. In other words, A is s -good if whenever y in (1.4) is of the form $y = Ax$ with s -sparse x , x is the unique optimal solution to (1.4).

Nullspace property. There is a simply-looking *necessary and sufficient* condition for a sensing matrix A to be s -good—the *nullspace property* originating from [70]. After this property is guessed, it is easy to see that it indeed is necessary and sufficient for s -goodness; we, however, prefer to *derive* this condition from the “first principles,” which can be easily done via Convex Optimization. Thus, in the case in question, as in many other cases, there is no necessity to be smart to arrive at the truth via a “lucky guess”; it suffices to be knowledgeable and use the standard tools.

Let us start with necessary condition for A to be such that whenever x is s -sparse, x is an optimal solution (perhaps not the unique one) of the optimization problem

$$\min_z \{\|z\|_1 : Az = Ax\}; \tag{P[x]}$$

we refer to the latter property of A as *weak s -goodness*. Our first observation is as follows:

Proposition 1.1. *If A is weakly s -good, then the following condition holds true: whenever I is a subset of $\{1, \dots, n\}$ of cardinality $\leq s$, we have*

$$\forall w \in \text{Ker}A \quad \|w_I\|_1 \leq \|w_{I^c}\|_1. \tag{1.6}$$

Proof is immediate. Assume A is weakly s -good, and let us verify (1.6). Let I be an s -element subset of $\{1, \dots, n\}$, and x be an s -sparse vector with support I . Since A is weakly s -good, x is an optimal solution to $(P[x])$. Rewriting the latter problem in the form of LP, that is, as

$$\min_{z,t} \left\{ \sum_j t_j : t_j + z_j \geq 0, t_j - z_j \geq 0, Az = Ax \right\},$$

and invoking LP optimality conditions, the necessary and sufficient condition for

⁵Note that in general x^s is not uniquely defined by x and s , since the s -th largest among the magnitudes of entries in x can be achieved at several entries. In our context, it does not matter how ties of this type are resolved; for the sake of definiteness, we can assume that when ordering the entries in x according to their magnitudes, from the largest to the smallest, entries of equal magnitude are ordered in the order of their indices.

$z = x$ to be the z -component of an optimal solution is the existence of λ_j^+ , λ_j^- , $\mu \in \mathbf{R}^m$ (Lagrange multipliers for the constraints $t_j - z_j \geq 0$, $t_j + z_j \geq 0$, and $Az = Ax$, respectively) such that

$$\begin{aligned}
 (a) \quad & \lambda_j^+ + \lambda_j^- = 1 \forall j, \\
 (b) \quad & \lambda^+ - \lambda^- + A^T \mu = 0, \\
 (c) \quad & \lambda_j^+ (|x_j| - x_j) = 0 \forall j, \\
 (d) \quad & \lambda_j^- (|x_j| + x_j) = 0 \forall j, \\
 (e) \quad & \lambda_j^+ \geq 0 \forall j, \\
 (f) \quad & \lambda_j^- \geq 0 \forall j.
 \end{aligned} \tag{1.7}$$

From (c, d), we have $\lambda_j^+ = 1, \lambda_j^- = 0$ for $j \in I_x^+$ and $\lambda_j^+ = 0, \lambda_j^- = 1$ for $j \in I_x^-$. From (a) and nonnegativity of λ_j^\pm it follows that for $j \notin I_x$ we should have $-1 \leq \lambda_j^+ - \lambda_j^- \leq 1$. With this in mind, the above optimality conditions admit eliminating λ 's and reduce to the following conclusion:

(!) x is an optimal solution to $(P[x])$ if and only if there exists vector $\mu \in \mathbf{R}^m$ such that the j -th entry of $A^T \mu$ is -1 if $x_j > 0$, $+1$ if $x_j < 0$, and a real from $[-1, 1]$ if $x_j = 0$.

Now let $w \in \text{Ker } A$ be a vector with the same signs of entries $w_i, i \in I$, as these of the entries in x . Then

$$\begin{aligned}
 0 = \mu^T A w &= [A^T \mu]^T w = \sum_j [A^T \mu]_j w_j \\
 \Rightarrow \sum_{j \in I_x} |w_j| &= \sum_{j \in I_x} [A^T \mu]_j w_j = - \sum_{j \notin I_x} [A^T \mu]_j w_j \leq \sum_{j \notin I_x} |w_j|
 \end{aligned}$$

(we have used the fact that $[A^T \mu]_j = \text{sign } x_j = \text{sign } w_j$ for $j \in I_x$ and $|[A^T \mu]_j| \leq 1$ for all j). Since I can be an arbitrary s -element subset of $\{1, \dots, n\}$ and the pattern of signs of an s -sparse vector x supported on I can be arbitrary, (1.6) holds true. \square

1.2.1.3 Nullspace property

In fact, it can be shown that (1.6) is not only a necessary, but also sufficient condition for weak s -goodness of A ; we, however, skip this verification, since our goal so far was to *guess* the condition for s -goodness, and this goal has already been achieved—from what we already know it immediately follows that a *necessary* condition for s -goodness is for the inequality in (1.6) to be strict whenever $w \in \text{Ker } A$ is nonzero. Indeed, we already know that if A is s -good, then for every I of cardinality s and every nonzero $w \in \text{Ker } A$ it holds

$$\|w_I\|_1 \leq \|w_{I^c}\|_1.$$

If the latter inequality for some I and w in question holds true as equality, then A clearly is *not* s -good, since the s -sparse signal $x = w_I$ is *not* the unique optimal solution to $(P[x])$ —the vector $-w_{I^c}$ is a different feasible solution to the same problem and with the same value of the objective. We conclude that for A to be s -good, a necessary condition is

$$\forall (0 \neq w \in \text{Ker } A, I, \text{Card}(I) \leq s) : \|w_I\|_1 < \|w_{I^c}\|_1.$$

By the standard compactness argument, this is the same as the existence of $\gamma \in (0, 1)$ such that

$$\forall (w \in \text{Ker } A, I, \text{Card}(I) \leq s) : \|w_I\|_1 \leq \gamma \|w_{I^c}\|_1,$$

or—which is the same—existence of $\kappa \in (0, 1/2)$ such that

$$\forall (w \in \text{Ker } A, I, \text{Card}(I) \leq s) : \|w_I\|_1 \leq \kappa \|w\|_1.$$

Finally, the supremum of $\|w_I\|_1$ over I of cardinality s is the norm $\|w\|_{s,1}$ (the sum of s largest magnitudes of entries) of w , so that the condition we are processing finally can be formulated as

$$\exists \kappa \in (0, 1/2) : \|w\|_{s,1} \leq \kappa \|w\|_1 \quad \forall w \in \text{Ker } A. \quad (1.8)$$

The resulting *nullspace condition* in fact is necessary and sufficient for A to be s -good:

Proposition 1.2. *Condition (1.8) is necessary and sufficient for A to be s -good.*

Proof. We have already seen that the nullspace condition is necessary for s -goodness. To verify sufficiency, let A satisfy the nullspace condition, and let us prove that A is s -good. Indeed, let x be an s -sparse vector, and y be an optimal solution to $(P[x])$; all we need is to prove that $y = x$. Let I be the support of x , and $w = y - x$, so that $w \in \text{Ker } A$. By the nullspace property we have

$$\begin{aligned} & \|w_I\|_1 \leq \kappa \|w\|_1 = \kappa (\|w_I\|_1 + \|w_{I^c}\|_1) = \kappa (\|w_I\|_1 + \|y_{I^c}\|_1) \\ \Rightarrow & \|w_I\|_1 \leq \frac{\kappa}{1-\kappa} \|y_{I^c}\|_1 \\ \Rightarrow & \|x\|_1 = \|x_I\|_1 = \|y_I - w_I\|_1 \leq \|y_I\|_1 + \frac{\kappa}{1-\kappa} \|y_{I^c}\|_1 \leq \|y_I\|_1 + \|y_{I^c}\|_1 = \|y\|_1 \end{aligned}$$

where the concluding \leq is due to $\kappa \in [0, 1/2)$. Since x is a feasible, and y is an optimal solution to $(P[x])$, the resulting inequality $\|x\|_1 \leq \|y\|_1$ must be equality, which, again due to $\kappa \in [0, 1/2)$, is possible only when $y_{I^c} = 0$. Thus, y has the same support I as x , and $w = x - y \in \text{Ker } A$ is supported on s -element set I ; by nullspace property, we should have $\|w_I\|_1 \leq \kappa \|w\|_1 = \kappa \|w_I\|_1$, which is possible only when $w = 0$. \square

1.2.2 Imperfect ℓ_1 minimization

We have found a necessary and sufficient condition for ℓ_1 minimization to recover *exactly* s -sparse signals in the *noiseless* case. More often than not, both these assumptions are violated: instead of s -sparse signals, we should speak about “nearly s -sparse” ones, quantifying the deviation from sparsity by the distance from the signal x underlying the observations to its best s -sparse approximation x^s . Similarly, we should allow for nonzero observation noise. With noisy observations and/or imperfect sparsity, we cannot hope to recover the signal exactly. All we may hope for, is to recover it with some error depending on the level of observation noise and “deviation from s -sparsity,” and tending to zero as the level and deviation tend to 0. We are about to quantify the nullspace property to allow for instructive “error analysis.”

1.2.2.1 Contrast matrices and quantifications of Nullspace property

By itself, the nullspace property says something about the signals from the kernel of the sensing matrix. We can reformulate it equivalently to say something important about *all* signals. Namely, observe that given sparsity s and $\kappa \in (0, 1/2)$, the nullspace property

$$\|w\|_{s,1} \leq \kappa \|w\|_1 \quad \forall w \in \text{Ker } A \quad (1.9)$$

is satisfied if and only if for a properly selected constant C one has⁶

$$\|w\|_{s,1} \leq C \|Aw\|_2 + \kappa \|w\|_1 \quad \forall w. \quad (1.10)$$

Indeed, (1.10) clearly implies (1.9); to get the inverse implication, note that for every h orthogonal to $\text{Ker } A$ it holds

$$\|Ah\|_2 \geq \sigma \|h\|_2,$$

where $\sigma > 0$ is the minimal positive singular value of A . Now, given $w \in \mathbf{R}^n$, we can decompose w into the sum of $\bar{w} \in \text{Ker } A$ and $h \in (\text{Ker } A)^\perp$, so that

$$\begin{aligned} \|w\|_{s,1} &\leq \|\bar{w}\|_{s,1} + \|h\|_{s,1} \leq \kappa \|\bar{w}\|_1 + \sqrt{s} \|h\|_{s,2} \leq \kappa (\|w\|_1 + \|h\|_1) + \sqrt{s} \|h\|_2 \\ &\leq \kappa \|w\|_1 + [\kappa\sqrt{n} + \sqrt{s}] \|h\|_2 \leq \underbrace{\sigma^{-1}[\kappa\sqrt{n} + \sqrt{s}]}_C \underbrace{\|Ah\|_2}_{=\|Aw\|_2} + \kappa \|w\|_1, \end{aligned}$$

as required in (1.10).

Condition $\mathbf{Q}_1(s, \kappa)$. For our purposes, it is convenient to present the condition (1.10) in the following flexible form:

$$\|w\|_{s,1} \leq s \|H^T Aw\| + \kappa \|w\|_1, \quad (1.11)$$

where H is an $m \times N$ *contrast* matrix and $\|\cdot\|$ is some norm on \mathbf{R}^N . Whenever a pair $(H, \|\cdot\|)$, called *contrast pair*, satisfies (1.11), we say that $(H, \|\cdot\|)$ *satisfies condition $\mathbf{Q}_1(s, \kappa)$* . From what we have seen, *If* A possesses nullspace property with some sparsity level s and some $\kappa \in (0, 1/2)$, *then* there are many ways to select pairs $(H, \|\cdot\|)$ satisfying $\mathbf{Q}_1(s, \kappa)$, e.g., to take $H = CI_m$ with appropriately large C and $\|\cdot\| = \|\cdot\|_2$.

Conditions $\mathbf{Q}_q(s, \kappa)$. As we will see in a while, it makes sense to embed the condition $\mathbf{Q}_1(s, \kappa)$ into a parametric family of conditions $\mathbf{Q}_q(s, \kappa)$, where the parameter q runs through $[1, \infty]$. Specifically,

Given an $m \times n$ sensing matrix A , sparsity level $s \leq n$, and $\kappa \in (0, 1/2)$, we say that $m \times N$ matrix H and a norm $\|\cdot\|$ on \mathbf{R}^N satisfy condition $\mathbf{Q}_q(s, \kappa)$ if

$$\|w\|_{s,q} \leq s^{\frac{1}{q}} \|H^T Aw\| + \kappa s^{\frac{1}{q}-1} \|w\|_1 \quad \forall w \in \mathbf{R}^n. \quad (1.12)$$

Let us make two immediate observations on relations between the conditions:

- A.** When a pair $(H, \|\cdot\|)$ satisfies condition $\mathbf{Q}_q(s, \kappa)$, the pair satisfies also all conditions $\mathbf{Q}_{q'}(s, \kappa)$ with $1 \leq q' \leq q$.

⁶Note that (1.9) is exactly the $\phi^2(s, \kappa)$ -Compatibility condition of [231] with $\phi(s, \kappa) = C/\sqrt{s}$; see also [232] for the analysis of relationships of this condition with other assumptions (e.g., a similar Restricted Eigenvalue assumption of [20]) used to analyse ℓ_1 -minimization procedures.

Indeed in the situation in question for $1 \leq q' \leq q$ it holds

$$\begin{aligned} \|w\|_{s,q'} &\leq s^{\frac{1}{q'} - \frac{1}{q}} \|w\|_{q,s} \leq s^{\frac{1}{q'} - \frac{1}{q}} \left[s^{\frac{1}{q}} \|H^T A w\| + \kappa s^{\frac{1}{q} - 1} \|w\|_1 \right] \\ &= s^{\frac{1}{q'}} \|H^T A w\| + \kappa s^{\frac{1}{q'} - 1} \|w\|_1, \end{aligned}$$

where the first inequality is the standard inequality between ℓ_p -norms of the s -dimensional vector w^s .

B. When a pair $(H, \|\cdot\|)$ satisfies condition $\mathbf{Q}_q(s, \kappa)$ and $1 \leq s' \leq s$, the pair $((s/s')^{\frac{1}{q}} H, \|\cdot\|)$ satisfies the condition $\mathbf{Q}_q(s', \kappa)$.

Indeed, in the situation in question we clearly have for $1 \leq s' \leq s$:

$$\|w\|_{s',q} \leq \|w\|_{s,q} \leq (s')^{\frac{1}{q}} \left\| \left[(s/s')^{\frac{1}{q}} H \right] A w \right\| + \kappa \underbrace{s^{\frac{1}{q} - 1}}_{\leq (s')^{\frac{1}{q} - 1}} \|w\|_1.$$

1.2.3 Regular ℓ_1 recovery

Given the observation scheme (1.1) with an $m \times n$ sensing matrix A , we define the *regular ℓ_1 recovery* of x via observation y as

$$\hat{x}_{\text{reg}}(y) \in \underset{u}{\text{Argmin}} \{ \|u\|_1 : \|H^T (Au - y)\| \leq \rho \}, \quad (1.13)$$

where the *contrast matrix* $H \in \mathbf{R}^{m \times N}$, the norm $\|\cdot\|$ on \mathbf{R}^N and $\rho > 0$ are parameters of the construction.

The role of \mathbf{Q} -conditions we have introduced is clear from the following

Theorem 1.3. *Let s be a positive integer, $q \in [1, \infty]$ and $\kappa \in (0, 1/2)$. Assume that a pair $(H, \|\cdot\|)$ satisfies the condition $\mathbf{Q}_q(s, \kappa)$ associated with A , and let*

$$\Xi_\rho = \{ \eta : \|H^T \eta\| \leq \rho \}. \quad (1.14)$$

Then for all $x \in \mathbf{R}^n$ and $\eta \in \Xi_\rho$ one has

$$\|\hat{x}_{\text{reg}}(Ax + \eta) - x\|_p \leq \frac{4(2s)^{\frac{1}{p}}}{1 - 2\kappa} \left[\rho + \frac{\|x - x^s\|_1}{2s} \right], \quad 1 \leq p \leq q. \quad (1.15)$$

The above result can be slightly strengthened by replacing the assumption that $(H, \|\cdot\|)$ satisfies $\mathbf{Q}_q(s, \kappa)$ with some $\kappa < 1/2$, with a weaker—by observation \mathbf{A} from Section 1.2.2.1—assumption that $(H, \|\cdot\|)$ satisfies $\mathbf{Q}_1(s, \varkappa)$ with $\varkappa < 1/2$ and satisfies $\mathbf{Q}_q(s, \kappa)$ with some (perhaps large) κ :

Theorem 1.4. *Given A , integer $s > 0$, and $q \in [1, \infty]$, assume that $(H, \|\cdot\|)$ satisfies the condition $\mathbf{Q}_1(s, \varkappa)$ with $\varkappa < 1/2$ and the condition $\mathbf{Q}_q(s, \kappa)$ with some $\kappa \geq \varkappa$, and let Ξ_ρ be given by (1.14). Then for all $x \in \mathbf{R}^n$ and $\eta \in \Xi_\rho$ it holds:*

$$\|\hat{x}_{\text{reg}}(Ax + \eta) - x\|_p \leq \frac{4(2s)^{\frac{1}{p}} [1 + \kappa - \varkappa]^{\frac{q(p-1)}{p(q-1)}}}{1 - 2\varkappa} \left[\rho + \frac{\|x - x^s\|_1}{2s} \right], \quad 1 \leq p \leq q. \quad (1.16)$$

For proofs of Theorems 1.3 and 1.4, see Section 1.5.1.

Before commenting on the above results, let us present their alternative versions.

1.2.4 Penalized ℓ_1 recovery

Penalized ℓ_1 recovery of signal x from its observation (1.1) is

$$\hat{x}_{\text{pen}}(y) \in \underset{u}{\text{Argmin}} \{ \|u\|_1 + \lambda \|H^T(Au - y)\| \}, \quad (1.17)$$

where $H \in \mathbf{R}^{m \times N}$, a norm $\|\cdot\|$ on \mathbf{R}^N , and a positive real λ are parameters of the construction.

Theorem 1.5. *Given A , positive integer s , and $q \in [1, \infty]$, assume that $(H, \|\cdot\|)$ satisfies the conditions $\mathbf{Q}_q(s, \kappa)$ and $\mathbf{Q}_1(s, \varkappa)$ with $\varkappa < 1/2$ and $\kappa \geq \varkappa$. Then*

(i) *Let $\lambda \geq 2s$. Then for all $x \in \mathbf{R}^n$, $y \in \mathbf{R}^m$ it holds:*

$$\|\hat{x}_{\text{pen}}(y) - x\|_p \leq \frac{4\lambda^{\frac{1}{p}}}{1-2\varkappa} \left[1 + \frac{\kappa\lambda}{2s} - \varkappa \right]^{\frac{q(p-1)}{p(q-1)}} \left[\|H^T(Ax - y)\| + \frac{\|x - x^s\|_1}{2s} \right], \quad 1 \leq p \leq q. \quad (1.18)$$

In particular, with $\lambda = 2s$ we have:

$$\|\hat{x}_{\text{pen}}(y) - x\|_p \leq \frac{4(2s)^{\frac{1}{p}}}{1-2\varkappa} \left[1 + \kappa - \varkappa \right]^{\frac{q(p-1)}{p(q-1)}} \left[\|H^T(Ax - y)\| + \frac{\|x - x^s\|_1}{2s} \right], \quad 1 \leq p \leq q. \quad (1.19)$$

(ii) *Let $\rho \geq 0$, and let Ξ_ρ be given by (1.14). Then for all $x \in \mathbf{R}^n$ and all $\eta \in \Xi_\rho$ one has:*

$$\begin{aligned} \lambda \geq 2s &\Rightarrow \\ \|\hat{x}_{\text{pen}}(Ax + \eta) - x\|_p &\leq \frac{4\lambda^{\frac{1}{p}}}{1-2\varkappa} \left[1 + \frac{\kappa\lambda}{2s} - \varkappa \right]^{\frac{q(p-1)}{p(q-1)}} \left[\rho + \frac{\|x - x^s\|_1}{2s} \right], \quad 1 \leq p \leq q; \\ \lambda = 2s &\Rightarrow \\ \|\hat{x}_{\text{pen}}(Ax + \eta) - x\|_p &\leq \frac{4(2s)^{\frac{1}{p}}}{1-2\varkappa} \left[1 + \kappa - \varkappa \right]^{\frac{q(p-1)}{p(q-1)}} \left[\rho + \frac{\|x - x^s\|_1}{2s} \right], \quad 1 \leq p \leq q. \end{aligned} \quad (1.20)$$

For proof, see Section 1.5.2.

1.2.5 Discussion

Some remarks are in order.

A. Qualitatively speaking, Theorems 1.3, 1.4, and 1.5 say the same thing: when **Q**-conditions are satisfied, the regular or penalized recoveries reproduce the true signal *exactly* when there is no observation noise and the signal is s -sparse. In the presence of observation error η and imperfect sparsity, the signal is recovered within the error which can be upper-bounded by the sum of two terms, one proportional to the magnitude of observation noise and one proportional to the deviation $\|x - x^s\|_1$ of the signal from s -sparse ones. In the penalized recovery, the observation error is measured in the scale given by the contrast matrix and the norm $\|\cdot\|$ —as $\|H^T\eta\|$ —and in the regular recovery by an a priori upper bound ρ on $\|H^T\eta\|$; when $\rho \geq \|H^T\eta\|$, η belongs to Ξ_ρ and thus the bounds (1.15) and (1.16) are applicable to the actual observation error η . Clearly, in qualitative terms, an error bound of this type is the best we may hope for. Now let us look at the quantitative aspect. Assume that in the regular recovery we use $\rho \approx \|H^T\eta\|$, and in the penalized one $\lambda = 2s$. In this case, error bounds (1.15), (1.16), and (1.20), up to factors C depending solely on \varkappa and κ , are the same, specifically,

$$\|\hat{x} - x\|_p \leq Cs^{1/p} [\|H^T\eta\| + \|x - x^s\|_1/s], \quad 1 \leq p \leq q. \quad (!)$$

Is this error bound bad or good? The answer depends on many factors, including on how well we select H and $\|\cdot\|$. To get a kind of orientation, consider the trivial case of *direct* observations, where matrix A is square and, moreover, is proportional to the unit matrix: $A = \alpha I$. Let us assume in addition that x is exactly s -sparse. In this case, the simplest way to ensure condition $\mathbf{Q}_q(s, \kappa)$, even with $\kappa = 0$, is to take $\|\cdot\| = \|\cdot\|_{s,q}$ and $H = s^{-1/q} \alpha^{-1} I$, so that (!) becomes

$$\|\hat{x} - x\|_p \leq C \alpha^{-1} s^{1/p-1/q} \|\eta\|_{s,q}, \quad 1 \leq p \leq q. \quad (!!)$$

As far as the dependence of the bound on the magnitude $\|\eta\|_{s,q}$ of the observation noise is concerned, this dependence is as good as it can be—even if we knew in advance the positions of the s entries of x of largest magnitudes, we would be unable to recover x in q -norm with error $\leq \alpha^{-1} \|\eta\|_{s,q}$. In addition, with the s largest magnitudes of entries in η equal to each other, the $\|\cdot\|_p$ -norm of the recovery error clearly cannot be guaranteed to be less than $\alpha^{-1} \|\eta\|_{s,p} = \alpha^{-1} s^{1/p-1/q} \|\eta\|_{s,q}$. Thus, at least for s -sparse signals x , our error bound is, basically, the best one can get already in the “ideal” case of direct observations.

B. Given that $(H, \|\cdot\|)$ obeys $\mathbf{Q}_1(s, \varkappa)$ with some $\varkappa < 1/2$, the larger the q such that the pair $(H, \|\cdot\|)$ obeys the condition $\mathbf{Q}_q(s, \kappa)$ with a given $\kappa \geq \varkappa$ (recall that κ can be $\geq 1/2$) and s , the larger the range $p \leq q$ of values of p where the error bounds (1.16) and (1.20) are applicable. This is in full accordance with the fact that if a pair $(H, \|\cdot\|)$ obeys condition $\mathbf{Q}_q(s, \kappa)$, it obeys also all conditions $\mathbf{Q}_{q'}(s, \kappa)$ with $1 \leq q' \leq q$ (item **A** in Section 1.2.2.1).

C. The flexibility offered by contrast matrix H and norm $\|\cdot\|$ allows us to adjust, to some extent, the recovery to the “geometry of observation errors.” For example, when η is “uncertain but bounded,” say, when all we know is that $\|\eta\|_2 \leq \delta$ with some given δ , all that matters (on the top of the requirement for $(H, \|\cdot\|)$ to obey \mathbf{Q} -conditions) is how large $\|H^T \eta\|$ could be when $\|\eta\|_2 \leq \delta$. In particular, when $\|\cdot\| = \|\cdot\|_2$, the error bound “is governed” by the spectral norm of H . Consequently, if we have a technique allowing us *to design* H such that $(H, \|\cdot\|_2)$ obeys \mathbf{Q} -condition(s) with given parameters, it makes sense to look for a design with as small a spectral norm of H as possible. In contrast to this, in the case of Gaussian noise the most interesting for applications,

$$y = Ax + \eta, \quad \eta \sim \mathcal{N}(0, \sigma^2 I_m), \quad (1.21)$$

looking at the spectral norm of H , with $\|\cdot\|_2$ in the role of $\|\cdot\|$, is counterproductive, since a typical realization of η is of Euclidean norm of order of $\sqrt{m}\sigma$ and thus is quite large when m is large. In this case to quantify “the magnitude” of $H^T \eta$ by the product of the spectral norm of H and the Euclidean norm of η is *completely misleading*—in typical cases, this product will grow rapidly with the number of observations m , completely ignoring the fact that η is random with zero mean.⁷ What is much better suited for the case of Gaussian noise, is the $\|\cdot\|_\infty$ norm in the role of $\|\cdot\|$ and the norm of H which is “the maximum of $\|\cdot\|_2$ -norms of the columns

⁷The simplest way to see the difference is to look at a particular entry $h^T \eta$ in $H^T \eta$. Operating with spectral norms, we upper-bound this entry by $\|h\|_2 \|\eta\|_2$, and the second factor for $\eta \sim \mathcal{N}(0, \sigma^2 I_m)$ is typically as large as $\sigma \sqrt{m}$. This is in sharp contrast to the fact that typical values of $h^T \eta$ are of order of $\sigma \|h\|_2$, independently of what m is!

in H ,” denoted by $\|H\|_{1,2}$. Indeed, with $\eta \sim \mathcal{N}(0, \sigma^2 I_m)$, the entries in $H^T \eta$ are Gaussian with zero mean and variance bounded by $\sigma^2 \|H\|_{1,2}^2$, so that $\|H^T \eta\|_\infty$ is the maximum of magnitudes of N zero mean Gaussian random variables with standard deviations bounded by $\sigma \|H\|_{1,2}$. As a result,

$$\text{Prob}\{\|H^T \eta\|_\infty \geq \rho\} \leq 2N \text{Erfc}\left(\frac{\rho}{\sigma \|H\|_{1,2}}\right) \leq N e^{-\frac{\rho^2}{2\sigma^2 \|H\|_{1,2}^2}}, \quad (1.22)$$

where

$$\text{Erfc}(s) = \text{Prob}_{\xi \sim \mathcal{N}(0,1)}\{\xi \geq s\} = \frac{1}{\sqrt{2\pi}} \int_s^\infty e^{-t^2/2} dt$$

is the (slightly rescaled) complementary error function.

It follows that the typical values of $\|H^T \eta\|_\infty$, $\eta \sim \mathcal{N}(0, \sigma^2 I_m)$ are of order of at most $\sigma \sqrt{\ln(N)} \|H\|_{1,2}$. In applications we consider in this chapter, we have $N = O(m)$, so that with σ and $\|H\|_{1,2}$ given, typical values $\|H^T \eta\|_\infty$ are nearly independent of m . The bottom line is that ℓ_1 minimization is capable of handling large-scale Gaussian observation noise incomparably better than “uncertain-but-bounded” observation noise of similar magnitude (measured in Euclidean norm).

D. As far as comparison of regular and penalized ℓ_1 recoveries with the same pair $(H, \|\cdot\|)$ is concerned, the situation is as follows. Assume for the sake of simplicity that $(H, \|\cdot\|)$ satisfies $\mathbf{Q}_q(s, \kappa)$ with some s and $\kappa < 1/2$, and let the observation error be random. Given $\epsilon \in (0, 1)$, let

$$\rho_\epsilon[H, \|\cdot\|] = \min\{\rho : \text{Prob}\{\eta : \|H^T \eta\| \leq \rho\} \geq 1 - \epsilon\}; \quad (1.23)$$

this is nothing but the smallest ρ such that

$$\text{Prob}\{\eta \in \Xi_\rho\} \geq 1 - \epsilon \quad (1.24)$$

(see (1.14)), and thus the smallest ρ for which the error bound (1.15) for the regular ℓ_1 recovery holds true with probability $1 - \epsilon$ (or at least the smallest ρ for which the latter claim is supported by Theorem 1.3). With $\rho = \rho_\epsilon[H, \|\cdot\|]$, the regular ℓ_1 recovery guarantees (and that is the best guarantee one can extract from Theorem 1.3) that

(#) For some set Ξ , $\text{Prob}\{\eta \in \Xi\} \geq 1 - \epsilon$, of “good” realizations of $\eta \sim \mathcal{N}(0, \sigma^2 I_m)$, one has

$$\|\widehat{x}(Ax + \eta) - x\|_p \leq \frac{4(2s)^{\frac{1}{p}}}{1 - 2\kappa} \left[\rho_\epsilon[H, \|\cdot\|] + \frac{\|x - x^s\|_1}{2s} \right], \quad 1 \leq p \leq q, \quad (1.25)$$

whenever $x \in \mathbf{R}^n$ and $\eta \in \Xi_\rho$.

The error bound (1.19) (where we set $\varkappa = \kappa$) says that (#) holds true for the penalized ℓ_1 recovery with $\lambda = 2s$. The latter observation suggests that the penalized ℓ_1 recovery associated with $(H, \|\cdot\|)$ and $\lambda = 2s$ is better than its regular counterpart, the reason being twofold. First, in order to ensure (#) with the regular recovery, the “built in” parameter ρ of this recovery should be set to $\rho_\epsilon[H, \|\cdot\|]$, and the latter quantity is not always easy to identify. In contrast to this, the construc-

tion of penalized ℓ_1 recovery is completely independent of a priori assumptions on the structure of observation errors, while automatically ensuring (#) for the error model we use. Second, and more importantly, for the penalized recovery the bound (1.25) is no more than the “worst, with confidence $1 - \epsilon$, case,” while the typical values of the quantity $\|H^T \eta\|$ which indeed participates in the error bound (1.18) may be essentially smaller than $\rho_\epsilon[H, \|\cdot\|]$. Numerical experience fully supports the above claim: the difference in observed performance of the two routines in question, although not dramatic, is definitely in favor of the penalized recovery. The only potential disadvantage of the latter routine is that the penalty parameter λ should be tuned to the level s of sparsity we aim at, while the regular recovery is free of any guess of this type. Of course, the “tuning” is rather loose—all we need (and experiments show that we indeed need this) is the relation $\lambda \geq 2s$, so that a rough upper bound on s will do. However, that bound (1.18) deteriorates as λ grows.

Finally, we remark that when H is $m \times N$ and $\eta \sim \mathcal{N}(0, \sigma^2 I_m)$, we have

$$\rho_\epsilon[H, \|\cdot\|_\infty] \leq \sigma \text{ErfcInv}\left(\frac{\epsilon}{2N}\right) \|H\|_{1,2} \leq \sigma \sqrt{2 \ln(N/\epsilon)} \|H\|_{1,2}$$

(see (1.22)); here $\text{ErfcInv}(\delta)$ is the inverse complementary error function:

$$\text{Erfc}(\text{ErfcInv}(\delta)) = \delta, \quad 0 < \delta < 1. \tag{1.26}$$

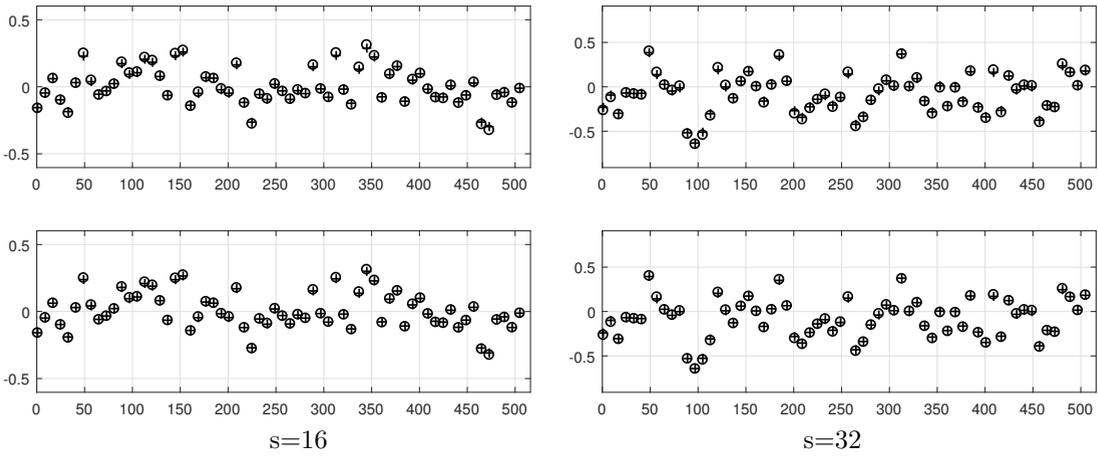
How it works. Here we present a small numerical illustration. We observe in Gaussian noise $m = n/2$ randomly selected terms in n -element “time series” $z = (z_1, \dots, z_n)$ and want to recover this series under the assumption that the series is “nearly s -sparse in frequency domain,” that is, that

$$z = Fx \text{ with } \|x - x^s\|_1 \leq \delta,$$

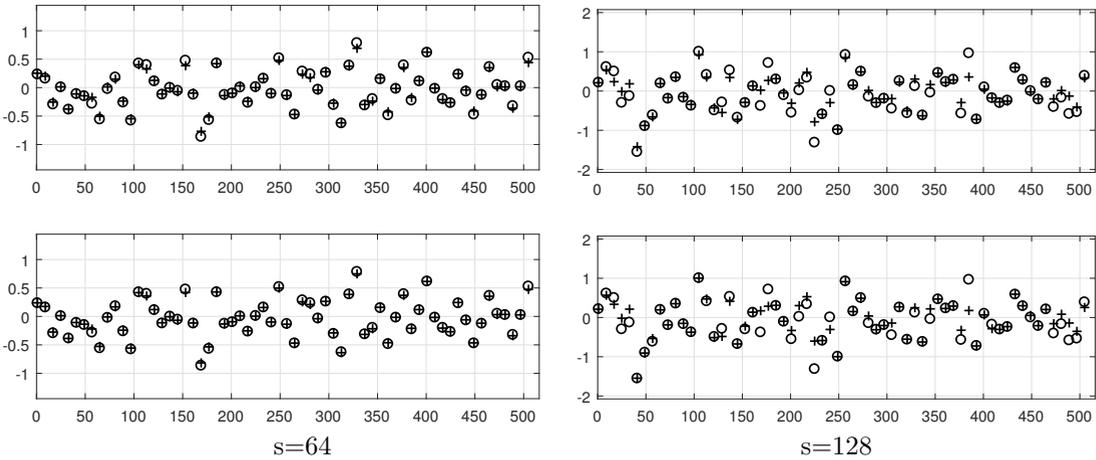
where F is the matrix of $n \times n$ the Inverse Discrete Cosine Transform, x^s is the vector obtained from x by zeroing out all but the s entries of largest magnitudes and δ upper-bounds the distance from x to s -sparse signals. Denoting by A the $m \times n$ submatrix of F corresponding to the time instants t where z_t is observed, our observation becomes

$$y = Ax + \sigma \xi,$$

where ξ is the standard Gaussian noise. After the signal in frequency domain, that is, x , is recovered by ℓ_1 minimization, let the recovery be \hat{x} , we recover the signal in the time domain as $\hat{z} = F\hat{x}$. In Figure 1.3, we present four test signals, of different (near-)sparsity, along with their regular and penalized ℓ_1 recoveries. The data in Figure 1.3 clearly show how the quality of ℓ_1 recovery deteriorates as the number s of “essential nonzeros” of the signal in the frequency domain grows. It is seen also that the penalized recovery meaningfully outperforms the regular one in the range of sparsities up to 64.



Top plots: regular ℓ_1 recovery, bottom plots: penalized ℓ_1 recovery.



Top plots: regular ℓ_1 recovery, bottom plots: penalized ℓ_1 recovery.

	$s = 16$	$s = 32$	$s = 64$	$s = 128$
$\ z - \hat{z}\ _2$	0.2417	0.3871	0.8178	4.8256
$\ z - \hat{z}\ _\infty$	0.0343	0.0514	0.1744	0.8272

recovery errors, regular ℓ_1 recovery

	$s = 16$	$s = 32$	$s = 64$	$s = 128$
$\ z - \hat{z}\ _2$	0.1399	0.2385	0.4216	5.3431
$\ z - \hat{z}\ _\infty$	0.0177	0.0362	0.1023	0.9141

recovery errors, penalized ℓ_1 recovery

Figure 1.3: Regular and penalized ℓ_1 recovery of nearly s -sparse signals. \circ : true signals, $+$: recoveries (to make the plots readable, one per eight consecutive vector's entries is shown). Problem sizes are $m = 256$ and $n = 2m = 512$, noise level is $\sigma = 0.01$, deviation from s -sparsity is $\|x - x^s\|_1 = 1$, contrast pair is $(H = \sqrt{n/m}A, \|\cdot\|_\infty)$. In penalized recovery, $\lambda = 2s$, parameter ρ of regular recovery is set to $\sigma \cdot \text{ErfcInv}(0.005/n)$.

1.3 VERIFIABILITY AND TRACTABILITY ISSUES

The good news about ℓ_1 recovery stated in Theorems 1.3, 1.4, and 1.5 is “conditional”—we assume that we are smart enough to point out a pair $(H, \|\cdot\|)$ satisfying condition $\mathbf{Q}_1(s, \varkappa)$ with $\varkappa < 1/2$ (and condition $\mathbf{Q}_q(s, \kappa)$ with a “moderate” \varkappa ⁸). The related issues are twofold:

1. First, we do not know in which range of s , m , and n these conditions, or even the weaker than $\mathbf{Q}_1(s, \varkappa)$, $\varkappa < 1/2$, nullspace property can be satisfied; and without the nullspace property, ℓ_1 minimization becomes useless, at least when we want to guarantee its validity whatever be the s -sparse signal we want to recover;
2. Second, it is unclear how to verify whether a given sensing matrix A satisfies the nullspace property for a given s , or a given pair $(H, \|\cdot\|)$ satisfies the condition $\mathbf{Q}_q(s, \kappa)$ with given parameters.

What is known about these crucial issues can be outlined as follows.

1. It is known that for given m, n with $m \ll n$ (say, $m/n \leq 1/2$), there exist $m \times n$ sensing matrices which are s -good for the values of s “nearly as large as m ,” specifically, for $s \leq O(1) \frac{m}{\ln(n/m)}$.⁹ Moreover, there are natural families of matrices where this level of goodness “is a rule.” E.g., when drawing an $m \times n$ matrix at random from Gaussian or Rademacher distributions (i.e., when filling the matrix with independent realizations of a random variable which is either a standard (zero mean, unit variance) Gaussian one, or takes values ± 1 with probabilities 0.5), the result will be s -good, for the outlined value of s , with probability approaching 1 as m and n grow. All this remains true when instead of speaking about matrices A satisfying “plain” nullspace properties, we are speaking about matrices A for which it is easy to point out a pair $(H, \|\cdot\|)$ satisfying the condition $\mathbf{Q}_2(s, \varkappa)$ with, say, $\varkappa = 1/4$.

The above results can be considered as a good news. A bad news is that we do *not* know how to check efficiently, given an s and a sensing matrix A , that the matrix is s -good, just as we do not know how to check that A admits good (i.e., satisfying $\mathbf{Q}_1(s, \varkappa)$ with $\varkappa < 1/2$) pairs $(H, \|\cdot\|)$. Even worse: we do not know an efficient recipe allowing us to build, given m , an $m \times 2m$ matrix A^m which is provably s -good for s larger than $O(1)\sqrt{m}$, which is a much smaller “level of goodness” than the one promised by theory for randomly generated matrices.¹⁰ The “common life” analogy of this situation would be as follows: you know that 90% of bricks in your wall are made of gold, and at the same time, you do not know how to tell a golden brick from a usual one.

2. There exist *verifiable sufficient conditions* for s -goodness of a sensing matrix, similarly to verifiable sufficient conditions for a pair $(H, \|\cdot\|)$ to satisfy condition

⁸ $\mathbf{Q}_q(s, \kappa)$ is always satisfied with “large enough” κ , e.g., $\kappa = s$, but such values of κ are of no interest: the associated bounds on p -norms of the recovery error are straightforward consequences of the bounds on the $\|\cdot\|_1$ -norm of this error yielded by the condition $\mathbf{Q}_1(s, \varkappa)$.

⁹Recall that $O(1)$ ’s denote positive *absolute constants*—appropriately chosen numbers like 0.5, or 1, or perhaps 100,000. We could, in principle, replace all $O(1)$ ’s with specific numbers; following the standard mathematical practice, we do not do it, partly out of laziness, partly because particular values of these numbers in our context are irrelevant.

¹⁰Note that the naive algorithm “generate $m \times 2m$ matrices at random until an s -good, with s promised by the theory, matrix is generated” is *not* an efficient recipe, since we still do not know how to check s -goodness efficiently.

$\mathbf{Q}_q(s, \kappa)$. The bad news is that when $m \ll n$, these verifiable sufficient conditions can be satisfied only when $s \leq O(1)\sqrt{m}$ —once again, in a much more narrow range of values of s than when typical randomly selected sensing matrices are s -good. In fact, $s = O(\sqrt{m})$ is so far *the best known* sparsity level for which we know individual s -good $m \times n$ sensing matrices with $m \leq n/2$.

1.3.1 Restricted Isometry Property and s -goodness of random matrices

There are several sufficient conditions for s -goodness, equally difficult to verify, but provably satisfied for typical random sensing matrices. The best known of them is the *Restricted Isometry Property* (RIP) defined as follows:

Definition 1.6. *Let k be an integer and $\delta \in (0, 1)$. We say that an $m \times n$ sensing matrix A possesses the Restricted Isometry Property with parameters δ and k , $\text{RIP}(\delta, k)$, if for every k -sparse $x \in \mathbf{R}^n$ one has*

$$(1 - \delta)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta)\|x\|_2^2. \quad (1.27)$$

It turns out that for natural ensembles of random $m \times n$ matrices, a typical matrix from the ensemble satisfies $\text{RIP}(\delta, k)$ with small δ and k “nearly as large as m ,” and that $\text{RIP}(\frac{1}{6}, 2s)$ implies the nullspace condition, and more. The simplest versions of the corresponding results are as follows.

Proposition 1.7. *Given $\delta \in (0, \frac{1}{5}]$, with properly selected positive $c = c(\delta)$, $d = d(\delta)$, $f = f(\delta)$ for all $m \leq n$ and all positive integers k such that*

$$k \leq \frac{m}{c \ln(n/m) + d} \quad (1.28)$$

the probability for a random $m \times n$ matrix A with independent $\mathcal{N}(0, \frac{1}{m})$ entries to satisfy $\text{RIP}(\delta, k)$ is at least $1 - \exp\{-fm\}$.

For proof, see Section 1.5.3.

Proposition 1.8. *Let $A \in \mathbf{R}^{m \times n}$ satisfy $\text{RIP}(\delta, 2s)$ for some $\delta < 1/3$ and positive integer s . Then*

(i) *The pair $(H = \frac{s^{-1/2}}{\sqrt{1-\delta}} I_m, \|\cdot\|_2)$ satisfies the condition $\mathbf{Q}_2\left(s, \frac{\delta}{1-\delta}\right)$ associated with A ;*

(ii) *The pair $(H = \frac{1}{1-\delta} A, \|\cdot\|_\infty)$ satisfies the condition $\mathbf{Q}_2\left(s, \frac{\delta}{1-\delta}\right)$ associated with A .*

For proof, see Section 1.5.4.

1.3.2 Verifiable sufficient conditions for $\mathbf{Q}_q(s, \kappa)$

When speaking about verifiable sufficient conditions for a pair $(H, \|\cdot\|)$ to satisfy $\mathbf{Q}_q(s, \kappa)$, it is convenient to restrict ourselves to the case where H , like A , is an $m \times n$ matrix, and $\|\cdot\| = \|\cdot\|_\infty$.

Proposition 1.9. *Let A be an $m \times n$ sensing matrix, and $s \leq n$ be a sparsity level.*

Given an $m \times n$ matrix H and $q \in [1, \infty]$, let us set

$$\nu_{s,q}[H] = \max_{j \leq n} \|\text{Col}_j[I - H^T A]\|_{s,q}, \quad (1.29)$$

where $\text{Col}_j[C]$ is j -th column of matrix C . Then

$$\|w\|_{s,q} \leq s^{1/q} \|H^T A w\|_\infty + \nu_{s,q}[H] \|w\|_1 \quad \forall w \in \mathbf{R}^n, \quad (1.30)$$

implying that the pair $(H, \|\cdot\|_\infty)$ satisfies the condition $\mathbf{Q}_q(s, s^{1-\frac{1}{q}} \nu_{s,q}[H])$.

Proof is immediate. Setting $V = I - H^T A$, we have

$$\begin{aligned} \|w\|_{s,q} &= \|[H^T A + V]w\|_{s,q} \leq \|H^T A w\|_{s,q} + \|Vw\|_{s,q} \\ &\leq s^{1/q} \|H^T A w\|_\infty + \sum_j |w_j| \|\text{Col}_j[V]\|_{s,q} \leq s^{1/q} \|H^T A\|_\infty + \nu_{s,q}[H] \|w\|_1. \quad \square \end{aligned}$$

Observe that the function $\nu_{s,q}[H]$ is an efficiently computable convex function of H , so that the set

$$\mathcal{H}_{s,q}^\kappa = \{H \in \mathbf{R}^{m \times n} : \nu_{s,q}[H] \leq s^{\frac{1}{q}-1} \kappa\} \quad (1.31)$$

is a computationally tractable convex set. When this set is nonempty for some $\kappa < 1/2$, every point H in this set is a contrast matrix such that $(H, \|\cdot\|_\infty)$ satisfies the condition $\mathbf{Q}_q(s, \kappa)$, that is, we can find contrast matrices making ℓ_1 minimization valid. Moreover, we can *design* contrast matrix, e.g., by minimizing over $\mathcal{H}_{s,q}^\kappa$ the function $\|H\|_{1,2}$, thus optimizing the sensitivity of the corresponding ℓ_1 recoveries to Gaussian observation noise; see items **C**, **D** in Section 1.2.5.

Explanation. The sufficient condition for s -goodness of A stated in Proposition 1.9 looks as if coming out of thin air; in fact it is a particular case of a simple and general construction as follows. Let $f(x)$ be a real-valued convex function on \mathbf{R}^n , and $X \subset \mathbf{R}^n$ be a nonempty bounded polytope represented as

$$X = \{x \in \text{Conv}\{g_1, \dots, g_N\} : Ax = 0\},$$

where $\text{Conv}\{g_1, \dots, g_N\} = \{\sum_i \lambda_i g_i : \lambda \geq 0, \sum_i \lambda_i = 1\}$ is the convex hull of vectors g_1, \dots, g_N . Our goal is to upper-bound the maximum $\text{Opt} = \max_{x \in X} f(x)$; this is a meaningful problem, since precisely maximizing a convex function over a polyhedron typically is a computationally intractable task. Let us act as follows: clearly, for any matrix H of the same size as A we have $\max_{x \in X} f(x) = \max_{x \in X} f([I - H^T A]x)$, since on X we have $[I - H^T A]x = x$. As a result,

$$\begin{aligned} \text{Opt} &:= \max_{x \in X} f(x) = \max_{x \in X} f([I - H^T A]x) \\ &\leq \max_{x \in \text{Conv}\{g_1, \dots, g_N\}} f([I - H^T A]x) \\ &= \max_{j \leq N} f([I - H^T A]g_j). \end{aligned}$$

We get a parametric—the parameter being H —upper bound on Opt , namely, the bound $\max_{j \leq N} f([I - H^T A]g_j)$. This parametric bound is convex in H , and thus is well suited for minimization over this parameter.

The result of Proposition 1.9 is inspired by this construction as applied to the

nullspace property: given an $m \times n$ sensing matrix A and setting

$$X = \{x \in \mathbf{R}^n : \|x\|_1 \leq 1, Ax = 0\} = \{x \in \text{Conv}\{\pm e_1, \dots, \pm e_n\} : Ax = 0\}$$

(e_i are the basic orths in \mathbf{R}^n), A is s -good if and only if

$$\text{Opt}_s := \max_{x \in X} \{f(x) := \|x\|_{s,1}\} < 1/2.$$

A verifiable sufficient condition for this, as yielded by the above construction, is the existence of an $m \times n$ matrix H such that

$$\max_{j \leq n} \max \{f([I_n - H^T A]e_j), f(-[I_n - H^T A]e_j)\} < 1/2,$$

or, which is the same,

$$\max_j \|\text{Col}_j[I_n - H^T A]\|_{s,1} < 1/2.$$

This observation brings to our attention the matrix $I - H^T A$ with varying H and the idea of expressing sufficient conditions for s -goodness and related properties in terms of this matrix.

1.3.3 Tractability of $\mathbf{Q}_\infty(s, \kappa)$

As we have already mentioned, the conditions $\mathbf{Q}_q(s, \kappa)$ are intractable, in the sense that we do not know how to verify whether a given pair $(H, \|\cdot\|)$ satisfies the condition. Surprisingly, this is *not* the case with the strongest of these conditions, the one with $q = \infty$. Namely,

Proposition 1.10. *Let A be an $m \times n$ sensing matrix, s be a sparsity level, and $\kappa \geq 0$. Then whenever a pair $(\bar{H}, \|\cdot\|)$ satisfies the condition $\mathbf{Q}_\infty(s, \kappa)$, there exists an $m \times n$ matrix H such that*

$$\|\text{Col}_j[I_n - H^T A]\|_{s,\infty} = \|\text{Col}_j[I_n - \bar{H}^T A]\|_\infty \leq s^{-1}\kappa, \quad 1 \leq j \leq n$$

(so that $(H, \|\cdot\|_\infty)$ satisfies $\mathbf{Q}_\infty(s, \kappa)$ by Proposition 1.9), and also

$$\|H^T \eta\|_\infty \leq \|\bar{H}^T \eta\| \quad \forall \eta \in \mathbf{R}^m. \quad (1.32)$$

In addition, the $m \times n$ contrast matrix H such that the pair $(H, \|\cdot\|_\infty)$ satisfies the condition $\mathbf{Q}_\infty(s, \kappa)$ with as small κ as possible can be found as follows. Consider n LP programs

$$\text{Opt}_i = \min_{\nu, h} \{\nu : \|A^T h - e^i\|_\infty \leq \nu\}, \quad (\#i)$$

where e^i is i -th basic orth of \mathbf{R}^n . Let $\text{Opt}_i, h_i, i = 1, \dots, n$ be optimal solutions to these problems; we set $H = [h_1, \dots, h_n]$; the corresponding value of κ is

$$\kappa_* = s \max_i \text{Opt}_i.$$

Besides this, there exists a transparent alternative description of the quantities Opt_i

(and thus of κ_*); specifically,

$$\text{Opt}_i = \max_x \{x_i : \|x\|_1 \leq 1, Ax = 0\}. \quad (1.33)$$

For proof, see Section 1.5.5.

Taken along with (1.32) and error bounds of Theorems 1.3, 1.4, and 1.5, Proposition 1.10 says that

As far as the condition $\mathbf{Q}_\infty(s, \kappa)$ is concerned, we lose nothing when restricting ourselves with pairs $(H \in \mathbf{R}^{m \times n}, \|\cdot\|_\infty)$ and contrast matrices H satisfying the condition

$$|[I_n - H^T A]_{ij}| \leq s^{-1} \kappa, \quad (1.34)$$

implying that $(H, \|\cdot\|_\infty)$ satisfies $\mathbf{Q}_\infty(s, \kappa)$.

The good news is that (1.34) is an explicit convex constraint on H (in fact, even on H and κ), so that we can solve the *design problems*, where we want to optimize a convex function of H under the requirement that $(H, \|\cdot\|_\infty)$ satisfies the condition $\mathbf{Q}_\infty(s, \kappa)$ (and, perhaps, additional convex constraints on H and κ).

1.3.3.1 Mutual Incoherence

The simplest (and up to some point in time, the only) verifiable sufficient condition for s -goodness of a sensing matrix A is expressed in terms of *mutual incoherence* of A , defined as

$$\mu(A) = \max_{i \neq j} \frac{|\text{Col}_i^T[A] \text{Col}_j[A]|}{\|\text{Col}_i[A]\|_2^2}. \quad (1.35)$$

This quantity is well defined whenever A has no zero columns (otherwise A is not even 1-good). Note that when A is normalized to have all columns of equal $\|\cdot\|_2$ -lengths,¹¹ $\mu(A)$ is small when the columns of A are nearly mutually orthogonal. The standard related result is that

Whenever A and a positive integer s are such that $\frac{2\mu(A)}{1+\mu(A)} < \frac{1}{s}$, A is s -good.

It is immediately seen that the latter condition is weaker than what we can get with the aid of (1.34):

Proposition 1.11. *Let A be an $m \times n$ matrix, and let the columns of $m \times n$ matrix H be given by*

$$\text{Col}_j(H) = \frac{1}{(1 + \mu(A)) \|\text{Col}_j(A)\|_2^2} \text{Col}_j(A), \quad 1 \leq j \leq n.$$

Then

$$|[I_m - H^T A]_{ij}| \leq \frac{\mu(A)}{1 + \mu(A)} \quad \forall i, j. \quad (1.36)$$

¹¹As far as ℓ_1 minimization is concerned, this normalization is non-restrictive: we always can enforce it by diagonal scaling of the signal underlying observations (1.1), and ℓ_1 minimization in scaled variables is the same as weighted ℓ_1 minimization in original variables.

In particular, when $\frac{2\mu(A)}{1+\mu(A)} < \frac{1}{s}$, A is s -good.

Proof. With H as above, the diagonal entries in $I - H^T A$ are equal to $1 - \frac{1}{1+\mu(A)} = \frac{\mu(A)}{1+\mu(A)}$, while by definition of mutual incoherence the magnitudes of the off-diagonal entries in $I - H^T A$ are $\leq \frac{\mu(A)}{1+\mu(A)}$ as well, implying (1.36). The “in particular” claim is given by (1.36) combined with Proposition 1.9. \square

1.3.3.2 From RIP to conditions $\mathbf{Q}_q(\cdot, \kappa)$

It turns out that when A is $\text{RIP}(\delta, k)$ and $q \geq 2$, it is easy to point out pairs $(H, \|\cdot\|)$ satisfying $\mathbf{Q}_q(t, \kappa)$ with a desired $\kappa > 0$ and properly selected t :

Proposition 1.12. *Let A be an $m \times n$ sensing matrix satisfying $\text{RIP}(\delta, 2s)$ with some s and some $\delta \in (0, 1)$, and let $q \in [2, \infty]$ and $\kappa > 0$ be given. Then*

(i) *Whenever a positive integer t satisfies*

$$t \leq \min \left[\left[\frac{\kappa(1-\delta)}{\delta} \right]^{\frac{q}{q-1}}, s^{\frac{q-2}{q-1}} \right] s^{\frac{q}{2q-2}}, \quad (1.37)$$

the pair $(H = \frac{t^{-\frac{1}{q}}}{\sqrt{1-\delta}} I_m, \|\cdot\|_2)$ satisfies $\mathbf{Q}_q(t, \kappa)$;

(ii) *Whenever a positive integer t satisfies (1.37), the pair $(H = \frac{s^{\frac{1}{2}} t^{-\frac{1}{q}}}{1-\delta} A, \|\cdot\|_\infty)$ satisfies $\mathbf{Q}_q(t, \kappa)$.*

For proof, see Section 1.5.4.

The most important consequence of Proposition 1.12 deals with the case of $q = \infty$ and states that *when s -goodness of a sensing matrix A can be ensured by the difficult to verify condition $\text{RIP}(\delta, 2s)$ with, say, $\delta = 0.2$, the somehow worse level of sparsity, $t = O(1)\sqrt{s}$ with properly selected absolute constant $O(1)$, can be certified via condition $\mathbf{Q}_\infty(t, \frac{1}{3})$ —there exists a pair $(H, \|\cdot\|_\infty)$ satisfying this condition. The point is that by Proposition 1.10, if the condition $\mathbf{Q}_\infty(t, \frac{1}{3})$ can at all be satisfied, a pair $(H, \|\cdot\|_\infty)$ satisfying this condition can be found efficiently.*

Unfortunately, the significant “dropdown” in the level of sparsity when passing from unverifiable RIP to verifiable \mathbf{Q}_∞ is inevitable; this bad news is what is on our agenda now.

1.3.3.3 Limits of performance of verifiable sufficient conditions for goodness

Proposition 1.13. *Let A be an $m \times n$ sensing matrix which is “essentially non-square,” specifically, such that $2m \leq n$, and let $q \in [1, \infty]$. Whenever a positive integer s and an $m \times n$ matrix H are linked by the relation*

$$\|\text{Col}_j[I_n - H^T A]\|_{s,q} < \frac{1}{2} s^{\frac{1}{q}-1}, \quad 1 \leq j \leq n, \quad (1.38)$$

one has

$$s \leq \sqrt{m}. \quad (1.39)$$

As a result, the sufficient condition for the validity of $\mathbf{Q}_q(s, \kappa)$ with $\kappa < 1/2$ from Proposition 1.9 can never be satisfied when $s > \sqrt{m}$. Similarly, the verifiable sufficient condition $\mathbf{Q}_\infty(s, \kappa)$, $\kappa < 1/2$, for s -goodness of A cannot be satisfied

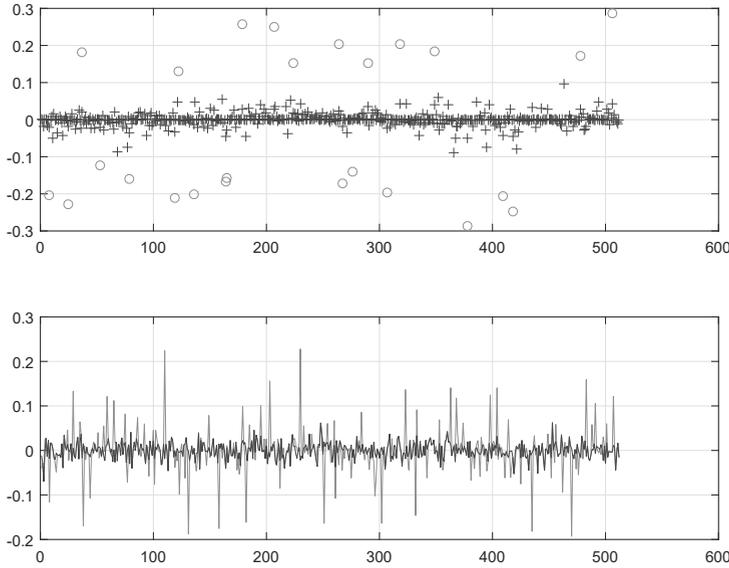


Figure 1.4: Erroneous ℓ_1 recovery of a 25-sparse signal, no observation noise. Top: frequency domain, o – true signal, + – recovery. Bottom: time domain.

when $s > \sqrt{m}$.

For proof, see Section 1.5.6.

We see that unless A is “nearly square,” our (same as all others known to us) verifiable sufficient conditions for s -goodness are unable to justify this property for “large” s . This unpleasant fact is in full accordance with the already mentioned fact that no individual provably s -good “essentially nonsquare” $m \times n$ matrices with $s \geq O(1)\sqrt{m}$ are known.

Matrices for which our verifiable sufficient conditions do establish s -goodness with $s \leq O(1)\sqrt{m}$ do exist.

How it works: Numerical illustration. Let us apply our machinery to the 256×512 randomly selected submatrix A of the matrix of 512×512 Inverse Discrete Cosine Transform which we used in experiments reported in Figure 1.3. These experiments exhibit nice performance of ℓ_1 minimization when recovering sparse (even nearly sparse) signals with as many as 64 nonzeros. *In fact, the level of goodness of A is at most 24,* as is witnessed in Figure 1.4.

In order to upper-bound the level of goodness of a matrix A , one can try to maximize the convex function $\|w\|_{s,1}$ over the set $W = \{w : Aw = 0, \|w\|_1 \leq 1\}$: if, for a given s , the maximum of $\|\cdot\|_{s,1}$ over W is $\geq 1/2$, the matrix is not s -good—it does not possess the nullspace property. Now, while global maximization of the convex function $\|w\|_{s,1}$ over W is difficult, we can try to find suboptimal solutions as follows. Let us start with a vector $w_1 \in W$ of $\|\cdot\|_1$ -norm 1, and let u^1 be obtained from w_1 by replacing the s entries in w_1 of largest magnitudes by the signs of these entries and zeroing out all other entries, so that $w_1^T u^1 = \|w_1\|_{s,1}$. After u^1 is found, let us solve the LO program $\max_w \{[u^1]^T w : w \in W\}$. w_1 is a feasible solution to this problem, so that for the optimal solution w_2 we have $[u^1]^T w_2 \geq [u^1]^T w_1 =$

$\|w_1\|_{s,1}$; this inequality, by virtue of what u^1 is, implies that $\|w_2\|_{s,1} \geq \|w_1\|_{s,1}$, and, by construction, $w_2 \in W$. We now can iterate the construction, with w_2 in the role of w_1 , to get $w_3 \in W$ with $\|w_3\|_{s,1} \geq \|w_2\|_{s,1}$, etc. Proceeding in this way, we generate a sequence of points from W with monotonically increasing value of the objective $\|\cdot\|_{s,1}$ we want to maximize. We terminate this recurrence either when the achieved value of the objective becomes $\geq 1/2$ (then we know for sure that A is not s -good, and can proceed to investigating s -goodness for a smaller value of s) or when the recurrence gets stuck—the observed progress in the objective falls below a given threshold, say, 10^{-6} . When it happens, we can restart the process from a new starting point randomly selected in W , after getting stuck, restart again, etc., until we exhaust our time budget. The output of the process is the best of the points we have generated—that of the largest $\|\cdot\|_{s,1}$. Applying this approach to the matrix A in question, in a couple of minutes it turns out that the matrix is at most 24-good.

One can ask how it may happen that previous experiments with recovering 64-sparse signals went fine, when in fact some 25-sparse signals cannot be recovered by ℓ_1 minimization even in the ideal noiseless case. The answer is simple: in our experiments, we dealt with *randomly selected* signals, and typical randomly selected data are much nicer, whatever be the purpose of a numerical experiment, than the worst-case data.

It is interesting to understand also which goodness we can certify using our verifiable sufficient conditions. Computations show that the fully verifiable (and strongest in our scale of sufficient conditions for s -goodness) condition $\mathbf{Q}_\infty(s, \varkappa)$ can be satisfied with $\varkappa < 1/2$ when s is as large as 7 and $\varkappa = 0.4887$, and *cannot* be satisfied with $\varkappa < 1/2$ when $s = 8$. As for Mutual Incoherence, it can only justify 3-goodness, no more. We can hardly be happy with the resulting bounds—goodness at least 7 and at most 24; however, it could be worse.

1.4 EXERCISES FOR CHAPTER 1

Exercise 1.1.

The k -th Hadamard matrix, \mathcal{H}_k (here k is a nonnegative integer), is the $n_k \times n_k$ matrix, $n_k = 2^k$, given by the recurrence

$$\mathcal{H}_0 = [1]; \mathcal{H}_{k+1} = \left[\begin{array}{c|c} \mathcal{H}_k & \mathcal{H}_k \\ \hline \mathcal{H}_k & -\mathcal{H}_k \end{array} \right].$$

In the sequel, we assume that $k > 0$. Now comes the exercise:

1. Check that \mathcal{H}_k is a symmetric matrix with entries ± 1 , and columns of the matrix are mutually orthogonal, so that $\mathcal{H}_k/\sqrt{n_k}$ is an orthogonal matrix.
2. Check that when $k > 0$, \mathcal{H}_k has just two distinct eigenvalues, $\sqrt{n_k}$ and $-\sqrt{n_k}$, each of multiplicity $m_k := 2^{k-1} = n_k/2$.
3. Prove that whenever f is an eigenvector of \mathcal{H}_k , one has

$$\|f\|_\infty \leq \|f\|_1/\sqrt{n_k}.$$

Derive from this observation the conclusion as follows:

Let $a_1, \dots, a_{m_k} \in \mathbf{R}^{n_k}$ be unit vectors orthogonal to each other which are eigenvectors of \mathcal{H}_k with eigenvalues $\sqrt{n_k}$ (by the above, the dimension of the eigenspace of \mathcal{H}_k associated with the eigenvalue $\sqrt{n_k}$ is m_k , so that the required a_1, \dots, a_{m_k} do exist), and let A be the $m_k \times n_k$ matrix with the rows $a_1^T, \dots, a_{m_k}^T$. For every $x \in \text{Ker } A$ it holds

$$\|x\|_\infty \leq \frac{1}{\sqrt{n_k}} \|x\|_1,$$

whence A satisfies the nullspace property whenever the sparsity s satisfies $2s < \sqrt{n_k} = \sqrt{2m_k}$. Moreover, there exists (and can be found efficiently) an $m_k \times n_k$ contrast matrix $H = H_k$ such that for every $s < \frac{1}{2}\sqrt{n_k}$, the pair $(H_k, \|\cdot\|_\infty)$ satisfies the condition $\mathbf{Q}_\infty(s, \kappa_s = \underbrace{s/\sqrt{n_k}}_{<1/2})$ associated

with A , and the $\|\cdot\|_2$ -norms of columns of H_k do not exceed $\sqrt{2\frac{\sqrt{n_k}+1}{\sqrt{n_k}}}$.

Note that the above conclusion yields a sequence of individual $(m_k = 2^{k-1}) \times (n_k = 2^k)$ sensing matrices, $k = 1, 2, \dots$, with “size ratio” $n_k/m_k = 2$, which make an efficiently verifiable condition for s -goodness, say, $\mathbf{Q}_\infty(s, \frac{1}{3})$, satisfiable in basically the entire range of values of s allowed by Proposition 1.13. It would be interesting to get similar “fully constructive” results for other size ratios, like $m : n = 1 : 4$, $m : n = 1 : 8$, etc.

Exercise 1.2.

[Follow-up to Exercise 1.1] Exercise 1.1 provides us with an explicitly given $(m = 512) \times (n = 1024)$ sensing matrix \bar{A} such that the efficiently verifiable condition $\mathbf{Q}_\infty(15, \frac{15}{32})$ is satisfiable; in particular, \bar{A} is 15-good. With all we know about limits of performance of verifiable sufficient conditions for goodness, how should we evaluate this specific sensing matrix? Could we point out a sensing matrix of the same size which is provably s -good for a value of s larger (or “much larger”) than 15?

We do not know the answer, and you are requested to explore some possibilities, including (but not reducing to—you are welcome to investigate more options!) the following ones.

1. Generate at random a sample of $m \times n$ sensing matrices A , compute their mutual incoherences, and look at how large are the goodness levels certified by these incoherences. What happens when the matrices are Gaussian (independent $\mathcal{N}(0, 1)$ entries) and Rademacher (independent entries taking values ± 1 with probabilities $1/2$)?
2. Generate at random a sample of $m \times n$ matrices with independent $\mathcal{N}(0, 1/m)$ entries. Proposition 1.7 suggests that a sample matrix A has good chances to satisfy $\text{RIP}(\delta, k)$ with some $\delta < 1/3$ and some k , and thus to be s -good (and even more than this, see Proposition 1.8) for every $s \leq k/2$. Of course, given A we cannot check whether the matrix indeed satisfies $\text{RIP}(\delta, k)$ with given δ, k ; what we can try to do is to certify that $\text{RIP}(\delta, k)$ does *not* take place. To this end, it suffices to select at random, say, 200 $m \times k$ submatrices \bar{A} of A and compute the eigenvalues of $\bar{A}^T \bar{A}$; if A possesses $\text{RIP}(\delta, k)$, all these eigenvalues should belong to the segment $[1 - \delta, 1 + \delta]$, and if in reality this does not happen, A definitely is not $\text{RIP}(\delta, k)$.

Exercise 1.3.

Let us start with a preamble. Consider a finite Abelian group; the only thing which matters for us is that such a group G is specified by a collection of $k \geq 1$ of positive integers ν_1, \dots, ν_k and is comprised of all collections $\omega = (\omega_1, \dots, \omega_k)$ where every ω_i is an integer from the range $\{0, 1, \dots, \nu_k - 1\}$; the group operation, denoted by \oplus , is

$$(\omega_1, \dots, \omega_k) \oplus (\omega'_1, \dots, \omega'_k) = ((\omega_1 + \omega'_1) \bmod \nu_1, \dots, (\omega_k + \omega'_k) \bmod \nu_k),$$

where $a \bmod b$ is the remainder, taking values in $\{0, 1, \dots, b - 1\}$, in the division of an integer a by a positive integer b , e.g., $5 \bmod 3 = 2$ and $6 \bmod 3 = 0$. Clearly, the cardinality of the above group G is $n_k = \nu_1 \nu_2 \dots \nu_k$. A *character* of group G is a homomorphism acting from G into the multiplicative group of complex numbers of modulus 1, or, in simple words, a complex-valued function $\chi(\omega)$ on G such that $|\chi(\omega)| = 1$ for all $\omega \in G$ and $\chi(\omega \oplus \omega') = \chi(\omega)\chi(\omega')$ for all $\omega, \omega' \in G$. Note that characters themselves form a group w.r.t. pointwise multiplication; clearly, all characters of our G are functions of the form

$$\chi((\omega_1, \dots, \omega_k)) = \mu_1^{\omega_1} \dots \mu_k^{\omega_k},$$

where μ_i are restricted to be roots of degree ν_i from 1: $\mu_i^{\nu_i} = 1$. It is immediately seen that the group G_* of characters of G is of the same cardinality $n_k = \nu_1 \dots \nu_k$ as G . We can associate with G the matrix \mathcal{F} of size $n_k \times n_k$; the columns in the matrix are indexed by the elements ω of G , the rows by the characters $\chi \in G_*$ of G , and the element in cell (χ, ω) is $\chi(\omega)$. The standard example here corresponds to $k = 1$, in which case \mathcal{F} clearly is the $\nu_1 \times \nu_1$ matrix of the Discrete Fourier Transform.

Now comes the exercise:

1. Verify that the above \mathcal{F} is, up to factor $\sqrt{n_k}$, a unitary matrix: denoting by \bar{a} the complex conjugate of a complex number a , $\sum_{\omega \in G} \chi(\omega) \chi'(\omega)$ is n_k or 0 depending on whether $\chi = \chi'$ or $\chi \neq \chi'$.
2. Let $\bar{\omega}, \bar{\omega}'$ be two elements of G . Prove that there exists a permutation Π of elements of G which maps $\bar{\omega}$ into $\bar{\omega}'$ and is such that

$$\text{Col}_{\Pi(\bar{\omega})}[\mathcal{F}] = D \text{Col}_{\bar{\omega}}[\mathcal{F}] \quad \forall \bar{\omega} \in G,$$

where D is diagonal matrix with diagonal entries $\chi(\bar{\omega}')/\chi(\bar{\omega})$, $\chi \in G_*$.

3. Consider the special case of the above construction where $\nu_1 = \nu_2 = \dots = \nu_k = 2$. Verify that in this case \mathcal{F} , up to permutation of rows and permutation of columns (these permutations depend on how we assign the elements of G and G_* their serial numbers), is exactly the Hadamard matrix \mathcal{H}_k .
4. Extract from the above the following fact: let m, k be positive integers such that $m \leq n_k := 2^k$, and let sensing matrix A be obtained from \mathcal{H}_k by selecting m distinct rows. Assume we want to find an $m \times n_k$ contrast matrix H such that the pair $(H, \|\cdot\|_\infty)$ satisfies the condition $\mathbf{Q}_\infty(s, \kappa)$ with as small a κ as possible; by Proposition 1.10, to this end we should solve n LP programs

$$\text{Opt}_i = \min_h \|e^i - A^T h\|_\infty,$$

where e^i is i -th basic orth in \mathbf{R}^n . Prove that with A coming from \mathcal{H}_k , all

these problems have the same optimal value, and optimal solutions to all of the problems are readily given by the optimal solution to just one of them.

Exercise 1.4.

Proposition 1.13 states that the verifiable condition $\mathbf{Q}_\infty(s, \kappa)$ can certify s -goodness of an “essentially nonsquare” (with $m \leq n/2$) $m \times n$ sensing matrix A only when s is small as compared to m , namely, $s \leq \sqrt{2m}$. The exercise to follow is aimed at investigating what happens when $m \times n$ “low” (with $m < n$) sensing matrix A is “nearly square”, meaning that $m^\circ = n - m$ is small as compared to n . Specifically, you should prove that for properly selected individual $(n - m^\circ) \times n$ matrices A the condition $\mathbf{Q}_\infty(s, \kappa)$ with $\kappa < 1/2$ is satisfiable when s is as large as $O(1)n/\sqrt{m^\circ}$.

1. Let $n = 2^k p$ with positive integer p and integer $k \geq 1$, and let $m^\circ = 2^{k-1}$. Given a $2m^\circ$ -dimensional vector u , let u^+ be an n -dimensional vector built as follows: we split indexes from $\{1, \dots, n = 2^k p\}$ into 2^k consecutive groups I_1, \dots, I_{2^k} , p elements per group, and all entries of u^+ with indexes from I_i are equal to the i -th entry, u_i , of vector u . Now let U be the linear subspace in \mathbf{R}^{2^k} comprised of all eigenvectors, with eigenvalue $\sqrt{2^k}$, of the Hadamard matrix \mathcal{H}_k —see Exercise 1.1—so that the dimension of U is $2^{k-1} = m^\circ$, and let L be given by

$$L = \{u^+ : u \in U\} \subset \mathbf{R}^n.$$

Clearly, L is a linear subspace in \mathbf{R}^n of dimension m° . Prove that

$$\forall x \in L : \|x\|_\infty \leq \frac{\sqrt{2m^\circ}}{n} \|x\|_1.$$

Conclude that if A is an $(n - m^\circ) \times n$ sensing matrix with $\text{Ker } A = L$, then the verifiable sufficient condition $\mathbf{Q}_\infty(s, \kappa)$ does certify s -goodness of A whenever

$$1 \leq s < \frac{n}{2\sqrt{2m^\circ}}.$$

2. Let L be an m° -dimensional subspace in \mathbf{R}^n . Prove that L contains a nonzero vector x with

$$\|x\|_\infty \geq \frac{\sqrt{m^\circ}}{n} \|x\|_1,$$

so that the condition $\mathbf{Q}_\infty(s, \kappa)$ cannot certify s -goodness of an $(n - m^\circ) \times n$ sensing matrix A whenever $s > O(1)n/\sqrt{m^\circ}$, for properly selected absolute constant $O(1)$.

Exercise 1.5.

Utilize the results of Exercise 1.3 in a numerical experiment as follows.

- select n as an integer power 2^k of 2, say, $n = 2^{10} = 1024$;
- select a “representative” sequence M of values of m , $1 \leq m < n$, including values of m close to n and “much smaller” than n , say,

$$M = \{2, 5, 8, 16, 32, 64, 128, 256, 512, 7, 896, 960, 992, 1008, 1016, 1020, 1022, 1023\};$$

- for every $m \in M$,

- generate at random an $m \times n$ submatrix A of the $n \times n$ Hadamard matrix \mathcal{H}_k and utilize the result of item 4 of Exercise 1.3 in order to find the largest s such that the s -goodness of A can be certified via the condition $\mathbf{Q}_\infty(\cdot, \cdot)$; call $s(m)$ the resulting value of s ;
- generate a moderate sample of Gaussian $m \times n$ sensing matrices A_i with independent $\mathcal{N}(0, 1/m)$ entries and use the construction from Exercise 1.2 to upper-bound the largest s for which a matrix from the sample satisfies $\text{RIP}(1/3, 2s)$; call $\widehat{s}(m)$ the largest—over your A_i 's—of the resulting upper bounds.

The goal of the exercise is to compare the computed values of $s(m)$ and $\widehat{s}(m)$; in other words, we again want to understand how “theoretically perfect” RIP compares to “conservative restricted scope” condition \mathbf{Q}_∞ .

1.5 PROOFS

1.5.1 Proofs of Theorem 1.3, 1.4

All we need is to prove Theorem 1.4, since Theorem 1.3 is the particular case $\varkappa = \kappa < 1/2$ of Theorem 1.4.

Let us fix $x \in \mathbf{R}^n$ and $\eta \in \Xi_\rho$, and let us set $\widehat{x} = \widehat{x}_{\text{reg}}(Ax + \eta)$. Let also $I \subset \{1, \dots, n\}$ be the set of indexes of the s entries in x of largest magnitudes, I^o be the complement of I in $\{1, \dots, n\}$, and, for $w \in \mathbf{R}^n$, w_I and w_{I^o} be the vectors obtained from w by zeroing entries with indexes $j \notin I$ and $j \notin I^o$, respectively, and keeping the remaining entries intact. Finally, let $z = \widehat{x} - x$.

1°. By the definition of Ξ_ρ and due to $\eta \in \Xi_\rho$, we have

$$\|H^T([Ax + \eta] - Ax)\| \leq \rho, \quad (1.40)$$

so that x is a feasible solution to the optimization problem specifying \widehat{x} , whence $\|\widehat{x}\|_1 \leq \|x\|_1$. We therefore have

$$\begin{aligned} \|\widehat{x}_{I^o}\|_1 &= \|\widehat{x}\|_1 - \|\widehat{x}_I\|_1 \leq \|x\|_1 - \|\widehat{x}_I\|_1 = \|x_I\|_1 + \|x_{I^o}\|_1 - \|\widehat{x}_I\|_1 \\ &\leq \|z_I\|_1 + \|x_{I^o}\|_1, \end{aligned} \quad (1.41)$$

and therefore

$$\|z_{I^o}\|_1 \leq \|\widehat{x}_{I^o}\|_1 + \|x_{I^o}\|_1 \leq \|z_I\|_1 + 2\|x_{I^o}\|_1.$$

It follows that

$$\|z\|_1 = \|z_I\|_1 + \|z_{I^o}\|_1 \leq 2\|z_I\|_1 + 2\|x_{I^o}\|_1. \quad (1.42)$$

Further, by definition of \widehat{x} we have $\|H^T([Ax + \eta] - A\widehat{x})\| \leq \rho$, which combines with (1.40) to imply that

$$\|H^T A(\widehat{x} - x)\| \leq 2\rho. \quad (1.43)$$

2°. Since $(H, \|\cdot\|)$ satisfies $\mathbf{Q}_1(s, \varkappa)$, we have

$$\|z\|_{s,1} \leq s\|H^T Az\| + \varkappa\|z\|_1.$$

By (1.43), it follows that $\|z\|_{s,1} \leq 2s\rho + \varkappa\|z\|_1$, which combines with the evident

inequality $\|z_I\| \leq \|z\|_{s,1}$ (recall that $\text{Card}(I) = s$) and with (1.42) to imply that

$$\|z_I\|_1 \leq 2s\rho + \varkappa\|z\|_1 \leq 2s\rho + 2\varkappa\|z_I\|_1 + 2\varkappa\|x_{I^c}\|_1,$$

whence

$$\|z_I\|_1 \leq \frac{2s\rho + 2\varkappa\|x_{I^c}\|_1}{1 - 2\varkappa}.$$

Invoking (1.42), we conclude that

$$\|z\|_1 \leq \frac{4s}{1 - 2\varkappa} \left[\rho + \frac{\|x_{I^c}\|_1}{2s} \right]. \quad (1.44)$$

3°. Since $(H, \|\cdot\|)$ satisfies $\mathbf{Q}_q(s, \kappa)$, we have

$$\|z\|_{s,q} \leq s^{\frac{1}{q}} \|H^T A z\| + \kappa s^{\frac{1}{q}-1} \|z\|_1,$$

which combines with (1.44) and (1.43) to imply that

$$\|z\|_{s,q} \leq s^{\frac{1}{q}} 2\rho + \kappa s^{\frac{1}{q}} \frac{4\rho + 2s^{-1}\|x_{I^c}\|_1}{1 - 2\varkappa} \leq \frac{4s^{\frac{1}{q}}[1 + \kappa - \varkappa]}{1 - 2\varkappa} \left[\rho + \frac{\|x_{I^c}\|_1}{2s} \right] \quad (1.45)$$

(we have taken into account that $\varkappa < 1/2$ and $\kappa \geq \varkappa$). Let θ be the $(s+1)$ -st largest magnitude of entries in z , and let $w = z - z^s$. Now (1.45) implies that

$$\theta \leq \|z\|_{s,q} s^{-\frac{1}{q}} \leq \frac{4[1 + \kappa - \varkappa]}{1 - 2\varkappa} \left[\rho + \frac{\|x_{I^c}\|_1}{2s} \right].$$

Hence invoking (1.44) we have

$$\begin{aligned} \|w\|_q &\leq \|w\|_\infty^{\frac{q-1}{q}} \|w\|_1^{\frac{1}{q}} \leq \theta^{\frac{q-1}{q}} \|z\|_1^{\frac{1}{q}} \\ &\leq \theta^{\frac{q-1}{q}} \frac{(4s)^{\frac{1}{q}}}{[1 - 2\varkappa]^{\frac{1}{q}}} \left[\rho + \frac{\|x_{I^c}\|_1}{2s} \right]^{\frac{1}{q}} \\ &\leq \frac{4s^{\frac{1}{q}}[1 + \kappa - \varkappa]^{\frac{q-1}{q}}}{1 - 2\varkappa} \left[\rho + \frac{\|x_{I^c}\|_1}{2s} \right]. \end{aligned}$$

Taking into account (1.45) and the fact that the supports of z^s and w do not intersect, we get

$$\begin{aligned} \|z\|_q &\leq 2^{\frac{1}{q}} \max[\|z^s\|_q, \|w\|_q] = 2^{\frac{1}{q}} \max[\|z\|_{s,q}, \|w\|_q] \\ &\leq \frac{4(2s)^{\frac{1}{q}}[1 + \kappa - \varkappa]}{1 - 2\varkappa} \left[\rho + \frac{\|x_{I^c}\|_1}{2s} \right]. \end{aligned}$$

This bound combines with (1.44), the Moment inequality,¹² and with the relation $\|x_{I^c}\|_1 = \|x - x^s\|_1$ to imply (1.16). \square

¹²The Moment inequality states that if (Ω, μ) is a space with measure and f is a μ -measurable real-valued function on Ω , then $\phi(\rho) = \ln \left(\int_\Omega |f(\omega)|^{\frac{1}{\rho}} \mu(d\omega) \right)^\rho$ is a convex function of ρ on every segment $\Delta \subset [0, 1]$ such that $\phi(\cdot)$ is well defined at the endpoints of Δ . As a corollary, when $x \in \mathbf{R}^n$ and $1 \leq p \leq q \leq \infty$, one has $\|x\|_p \leq \|x\|_1^{\frac{q-p}{p(q-1)}} \|x\|_q^{\frac{q(p-1)}{p(q-1)}}$.

1.5.2 Proof of Theorem 1.5

Let us prove (i). Let us fix $x \in \mathbf{R}^n$ and η , and let us set $\hat{x} = \hat{x}_{\text{pen}}(Ax + \eta)$. Let also $I \subset \{1, \dots, K\}$ be the set of indexes of the s entries in x of largest magnitudes, I^o be the complement of I in $\{1, \dots, n\}$, and, for $w \in \mathbf{R}^n$, w_I and w_{I^o} be the vectors obtained from w by zeroing out all entries with indexes not in I and not in I^o , respectively. Finally, let $z = \hat{x} - x$ and $\nu = \|H^T \eta\|$.

1°. We have

$$\|\hat{x}\|_1 + \lambda \|H^T(A\hat{x} - Ax - \eta)\| \leq \|x\|_1 + \lambda \|H^T \eta\|$$

and

$$\|H^T(A\hat{x} - Ax - \eta)\| = \|H^T(Az - \eta)\| \geq \|H^T Az\| - \|H^T \eta\|,$$

whence

$$\|\hat{x}\|_1 + \lambda \|H^T Az\| \leq \|x\|_1 + 2\lambda \|H^T \eta\| = \|x\|_1 + 2\lambda \nu. \quad (1.46)$$

We have

$$\begin{aligned} \|\hat{x}\|_1 &= \|x + z\|_1 = \|x_I + z_I\|_1 + \|x_{I^o} + z_{I^o}\|_1 \\ &\geq \|x_I\|_1 - \|z_I\|_1 + \|z_{I^o}\|_1 - \|x_{I^o}\|_1, \end{aligned}$$

which combines with (1.46) to imply that

$$\|x_I\|_1 - \|z_I\|_1 + \|z_{I^o}\|_1 - \|x_{I^o}\|_1 + \lambda \|H^T Az\| \leq \|x\|_1 + 2\lambda \nu,$$

or, which is the same,

$$\|z_{I^o}\|_1 - \|z_I\|_1 + \lambda \|H^T Az\| \leq 2\|x_{I^o}\|_1 + 2\lambda \nu. \quad (1.47)$$

Since $(H, \|\cdot\|)$ satisfies $\mathbf{Q}_1(s, \varkappa)$, we have

$$\|z_I\|_1 \leq \|z\|_{s,1} \leq s \|H^T Az\| + \varkappa \|z\|_1,$$

so that

$$(1 - \varkappa) \|z_I\|_1 - \varkappa \|z_{I^o}\|_1 - s \|H^T Az\| \leq 0. \quad (1.48)$$

Taking a weighted sum of (1.47) and (1.48), the weights being 1 and 2, respectively, we get

$$(1 - 2\varkappa) [\|z_I\|_1 + \|z_{I^o}\|_1] + (\lambda - 2s) \|H^T Az\| \leq 2\|x_{I^o}\|_1 + 2\lambda \nu,$$

whence, due to $\lambda \geq 2s$,

$$\|z\|_1 \leq \frac{2\lambda \nu + 2\|x_{I^o}\|_1}{1 - 2\varkappa} \leq \frac{2\lambda}{1 - 2\varkappa} \left[\nu + \frac{\|x_{I^o}\|_1}{2s} \right]. \quad (1.49)$$

Further, by (1.46) we have

$$\lambda \|H^T Az\| \leq \|x\|_1 - \|\hat{x}\|_1 + 2\lambda \nu \leq \|z\|_1 + 2\lambda \nu,$$

which combines with (1.49) to imply that

$$\lambda \|H A^T z\| \leq \frac{2\lambda \nu + 2\|x_{I^o}\|_1}{1 - 2\varkappa} + 2\lambda \nu = \frac{2\lambda \nu(2 - 2\varkappa) + 2\|x_{I^o}\|_1}{1 - 2\varkappa}. \quad (1.50)$$

From $\mathbf{Q}_q(s, \kappa)$ it follows that

$$\|z\|_{s,q} \leq s^{\frac{1}{q}} \|H^T A z\| + \kappa s^{\frac{1}{q}-1} \|z\|_1,$$

which combines with (1.50) and (1.49) to imply that

$$\begin{aligned} \|z\|_{s,q} &\leq s^{\frac{1}{q}-1} [s \|H^T A z\| + \kappa \|z\|_1] \leq s^{\frac{1}{q}-1} \left[\frac{4s\nu(1-\varkappa) + \frac{2s}{\lambda} \|x_{I^c}\|_1}{1-2\varkappa} + \frac{\kappa[2\lambda\nu + \frac{\lambda}{s} \|x_{I^c}\|_1]}{1-2\varkappa} \right] \\ &= s^{\frac{1}{q}} \frac{[4(1-\varkappa) + 2s^{-1}\lambda\kappa]\nu + [2\lambda^{-1} + \kappa s^{-2}\lambda] \|x_{I^c}\|_1}{1-2\varkappa} \leq 4 \frac{s^{\frac{1}{q}}}{1-2\varkappa} [1 + \frac{\kappa\lambda}{2s} - \varkappa] \left[\nu + \frac{\|x_{I^c}\|_1}{2s} \right] \end{aligned} \quad (1.51)$$

(recall that $\lambda \geq 2s$, $\kappa \geq \varkappa$, and $\varkappa < 1/2$). It remains to repeat the reasoning following (1.45) in item 3^o of the proof of Theorem 1.4. Specifically, denoting by θ the $(s+1)$ -st largest magnitude of entries in z , (1.51) implies that

$$\theta \leq s^{-1/q} \|z\|_{s,q} \leq \frac{4}{1-2\varkappa} [1 + \kappa \frac{\lambda}{2s} - \varkappa] \left[\nu + \frac{\|x_{I^c}\|_1}{2s} \right], \quad (1.52)$$

so that for the vector $w = z - z^s$ one has

$$\|w\|_q \leq \theta^{1-\frac{1}{q}} \|w\|_1^{\frac{1}{q}} \leq \frac{4(\lambda/2)^{\frac{1}{q}}}{1-2\varkappa} [1 + \kappa \frac{\lambda}{2s} - \varkappa]^{\frac{q-1}{q}} \left[\nu + \frac{\|x_{I^c}\|_1}{2s} \right]$$

(we have used (1.52) and (1.49)). Hence, taking into account that z^s and w have nonintersecting supports,

$$\begin{aligned} \|z\|_q &\leq 2^{\frac{1}{q}} \max[\|z^s\|_q, \|w\|_q] = 2^{\frac{1}{q}} \max[\|z\|_{s,q}, \|w\|_q] \\ &\leq \frac{4\lambda^{\frac{1}{q}}}{1-2\varkappa} [1 + \kappa \frac{\lambda}{2s} - \varkappa] \left[\nu + \frac{\|x_{I^c}\|_1}{2s} \right] \end{aligned}$$

(we have used (1.51) along with $\lambda \geq 2s$ and $\kappa \geq \varkappa$). This combines with (1.49) and the Moment inequality to imply (1.18). All remaining claims of Theorem 1.5 are immediate corollaries of (1.18). \square

1.5.3 Proof of Proposition 1.7

1^o. Assuming $k \leq m$ and selecting a set I of k indices from $\{1, \dots, n\}$ distinct from each other, consider an $m \times k$ submatrix A_I of A comprised of columns with indexes from I , and let u be a unit vector in \mathbf{R}^k . The entries in the vector $m^{1/2} A_I u$ are independent $\mathcal{N}(0, 1)$ random variables, so that for the random variable $\zeta_u = \sum_{i=1}^m (m^{1/2} A_I u)_i^2$ and $\gamma \in (-1/2, 1/2)$ it holds (in what follows, expectations and probabilities are taken w.r.t. our ensemble of random A 's)

$$\ln(\mathbf{E}\{\exp\{\gamma\zeta\}\}) = m \ln \left(\frac{1}{\sqrt{2\pi}} \int e^{\gamma t^2 - \frac{1}{2}t^2} ds \right) = -\frac{m}{2} \ln(1 - 2\gamma).$$

Given $\alpha \in (0, 0.1]$ and selecting γ in such a way that $1 - 2\gamma = \frac{1}{1+\alpha}$, we get $0 < \gamma < 1/2$ and therefore

$$\begin{aligned} \text{Prob}\{\zeta_u > m(1 + \alpha)\} &\leq \mathbf{E}\{\exp\{\gamma\zeta_u\}\} \exp\{-m\gamma(1 + \alpha)\} \\ &= \exp\left\{-\frac{m}{2} \ln(1 - 2\gamma) - m\gamma(1 + \alpha)\right\} \\ &= \exp\left\{\frac{m}{2} [\ln(1 + \alpha) - \alpha]\right\} \leq \exp\left\{-\frac{m}{5}\alpha^2\right\}, \end{aligned}$$

and similarly, selecting γ in such a way that $1 - 2\gamma = \frac{1}{1-\alpha}$, we get $-1/2 < \gamma < 0$ and therefore

$$\begin{aligned} \text{Prob}\{\zeta_u < m(1 - \alpha)\} &\leq \mathbf{E}\{\exp\{\gamma\zeta_u\}\} \exp\{-m\gamma(1 - \alpha)\} \\ &= \exp\left\{-\frac{m}{2} \ln(1 - 2\gamma) - m\gamma(1 - \alpha)\right\} \\ &= \exp\left\{\frac{m}{2} [\ln(1 - \alpha) + \alpha]\right\} \leq \exp\left\{-\frac{m}{5}\alpha^2\right\}, \end{aligned}$$

and we end up with

$$u \in \mathbf{R}^k, \|u\|_2 = 1 \Rightarrow \begin{cases} \text{Prob}\{A : \|A_I u\|_2^2 > 1 + \alpha\} \leq \exp\left\{-\frac{m}{5}\alpha^2\right\} \\ \text{Prob}\{A : \|A_I u\|_2^2 < 1 - \alpha\} \leq \exp\left\{-\frac{m}{5}\alpha^2\right\} \end{cases}. \quad (1.53)$$

2°. As above, let $\alpha \in (0, 0.1]$, let

$$M = 1 + 2\alpha, \epsilon = \frac{\alpha}{2(1 + 2\alpha)},$$

and let us build an ϵ -net on the unit sphere S in \mathbf{R}^k as follows. We start with a point $u_1 \in S$; after $\{u_1, \dots, u_t\} \subset S$ is already built, we check whether there is a point in S at the $\|\cdot\|_2$ -distance from all points of the set $> \epsilon$. If it is the case, we add such a point to the net built so far and proceed with building the net; otherwise we terminate with the net $\{u_1, \dots, u_t\}$. By compactness of S and due to $\epsilon > 0$, this process eventually terminates; upon termination, we have at our disposal the collection $\{u_1, \dots, u_N\}$ of unit vectors such that every two of them are at $\|\cdot\|_2$ -distance $> \epsilon$ from each other, and every point from S is at distance at most ϵ from some point of the collection. We claim that the cardinality N of the resulting set can be bounded as

$$N \leq \left[\frac{2 + \epsilon}{\epsilon}\right]^k = \left[\frac{4 + 9\alpha}{\alpha}\right]^k \leq \left(\frac{5}{\alpha}\right)^k. \quad (1.54)$$

Indeed, the interiors of the $\|\cdot\|_2$ -balls of radius $\epsilon/2$ centered at the points u_1, \dots, u_N are mutually disjoint, and their union is contained in the $\|\cdot\|_2$ -ball of radius $1 + \epsilon/2$ centered at the origin; comparing the volume of the union and that of the ball, we arrive at (1.54).

3°. Consider event E comprised of all realizations of A such that for all k -element subsets I of $\{1, \dots, n\}$ and all $t \leq n$ it holds

$$1 - \alpha \leq \|A_I u_t\|_2^2 \leq 1 + \alpha. \quad (1.55)$$

By (1.53) and the union bound,

$$\text{Prob}\{A \notin E\} \leq 2N \binom{n}{k} \exp\left\{-\frac{m}{5}\alpha^2\right\}. \quad (1.56)$$

We claim that

$$A \in E \Rightarrow (1 - 2\alpha) \leq \|A_I u\|_2^2 \leq 1 + 2\alpha \quad \forall \left(\begin{array}{l} I \subset \{1, \dots, n\} : \text{Card}(I) = k \\ u \in \mathbf{R}^k : \|u\|_2 = 1 \end{array} \right). \quad (1.57)$$

Indeed, let $A \in E$, let us fix $I \in \{1, \dots, n\}$, $\text{Card}(I) = k$, and let M be the maximal value of the quadratic form $f(u) = u^T A_I^T A_I u$ on the unit $\|\cdot\|_2$ -ball B , centered at the origin, in \mathbf{R}^k . In this ball, f is Lipschitz continuous with constant $2M$ w.r.t. $\|\cdot\|_2$; denoting by \bar{u} a maximizer of the form on B , we lose nothing when assuming that \bar{u} is a unit vector. Now let u_s be the point of our net which is at $\|\cdot\|_2$ -distance at most ϵ from \bar{u} . We have

$$M = f(\bar{u}) \leq f(u_s) + 2M\epsilon \leq 1 + \alpha + 2M\epsilon,$$

whence

$$M \leq \frac{1 + \alpha}{1 - 2\epsilon} = 1 + 2\alpha,$$

implying the right inequality in (1.57). Now let u be unit vector in \mathbf{R}^k , and u_s be a point in the net at $\|\cdot\|$ -distance $\leq \epsilon$ from u . We have

$$f(u) \geq f(u_s) - 2M\epsilon \geq 1 - \alpha - 2\frac{1 + \alpha}{1 - 2\epsilon}\epsilon = 1 - 2\alpha,$$

justifying the first inequality in (1.57).

The bottom line is:

$$\begin{aligned} & \delta \in (0, 0.2], 1 \leq k \leq n \\ \Rightarrow \text{Prob}\{A : A \text{ does not satisfy RIP}(\delta, k)\} & \leq 2 \underbrace{\left(\frac{10}{\delta}\right)^k}_{\leq \left(\frac{20}{\delta}\right)^k} \binom{n}{k} \exp\left\{-\frac{m\delta^2}{20}\right\}. \end{aligned} \quad (1.58)$$

Indeed, setting $\alpha = \delta/2$, we have seen that whenever $A \notin E$, we have $(1 - \delta) \leq \|Au\|_2^2 \leq (1 + \delta)$ for all unit k -sparse u , which is nothing but $\text{RIP}(\delta, k)$; with this in mind, (1.58) follows from (1.56) and (1.54).

4°. It remains to verify that with properly selected—depending solely on δ —positive quantities c, d, f , for every $k \geq 1$ satisfying (1.28) the right-hand side in (1.58) is at most $\exp\{-fm\}$. Passing to logarithms, our goal is to ensure the relation

$$\begin{aligned} G := a(\delta)m - b(\delta)k - \ln \binom{n}{k} & \geq mf(\delta) > 0 \\ \left[a(\delta) = \frac{\delta^2}{20}, b(\delta) = \ln \left(\frac{20}{\delta}\right) \right] & \end{aligned} \quad (1.59)$$

provided that $k \geq 1$ satisfies (1.28).

Let k satisfy (1.28) with some c, d to be specified later, and let $y = k/m$. Assuming $d \geq 3$, we have $0 \leq y \leq 1/3$. Now, it is well known that

$$C := \ln \binom{n}{k} \leq n \left[\frac{k}{n} \ln \left(\frac{n}{k}\right) + \frac{n-k}{n} \ln \left(\frac{n}{n-k}\right) \right],$$

whence

$$\begin{aligned} C & \leq n \left[\frac{m}{n} y \ln \left(\frac{n}{my}\right) + \frac{n-k}{n} \underbrace{\ln \left(1 + \frac{k}{n-k}\right)}_{\leq \frac{k}{n-k}} \right] \\ & \leq n \left[\frac{m}{n} y \ln \left(\frac{n}{my}\right) + \frac{k}{n} \right] = m \left[y \ln \left(\frac{n}{my}\right) + y \right] \leq 2my \ln \left(\frac{n}{my}\right) \end{aligned}$$

(recall that $n \geq m$ and $y \leq 1/3$). It follows that

$$\begin{aligned} G &= a(\delta)m - b(\delta)k - C \geq a(\delta)m - b(\delta)ym - 2my \ln\left(\frac{n}{my}\right) \\ &= m \underbrace{\left[a(\delta) - b(\delta)y - 2y \ln\left(\frac{n}{m}\right) - 2y \ln\left(\frac{1}{y}\right) \right]}_H, \end{aligned}$$

and all we need is to select c, d in such a way that (1.28) would imply that $H \geq f$ with some positive $f = f(\delta)$. This is immediate: we can find $u(\delta) > 0$ such that when $0 \leq y \leq u(\delta)$, we have $2y \ln(1/y) + b(\delta)y \leq \frac{1}{3}a(\delta)$; selecting $d(\delta) \geq 3$ large enough, (1.28) would imply $y \leq u(\delta)$, and thus would imply

$$H \geq \frac{2}{3}a(\delta) - 2y \ln\left(\frac{n}{m}\right).$$

Now we can select $c(\delta)$ large enough for (1.28) to ensure that $2y \ln\left(\frac{n}{m}\right) \leq \frac{1}{3}a(\delta)$. With the c, d just specified, (1.28) implies that $H \geq \frac{1}{3}a(\delta)$, and we can take the latter quantity as $f(\delta)$. \square

1.5.4 Proof of Propositions 1.8 and 1.12

Let $x \in \mathbf{R}^n$, and let x^1, \dots, x^q be obtained from x by the following construction: x^1 is obtained from x by zeroing all but the s entries of largest magnitudes; x^2 is obtained by the same procedure applied to $x - x^1$; x^3 —by the same procedure applied to $x - x^1 - x^2$; and so on; the process is terminated at the first step q when it happens that $x = x^1 + \dots + x^q$. Note that for $j \geq 2$ we have $\|x^j\|_\infty \leq s^{-1}\|x^{j-1}\|_1$ and $\|x^j\|_1 \leq \|x^{j-1}\|_1$, whence also $\|x^j\|_2 \leq \sqrt{\|x^j\|_\infty \|x^j\|_1} \leq s^{-1/2}\|x^{j-1}\|_1$. It is easily seen that if A is RIP($\delta, 2s$), then for every two s -sparse vectors u, v with nonoverlapping supports we have

$$|v^T A^T A u| \leq \delta \|u\|_2 \|v\|_2. \quad (*)$$

Indeed, for s -sparse u, v , let I be the index set of cardinality $\leq 2s$ containing the supports of u and v , so that, denoting by A_I the submatrix of A comprised of columns with indexes from I , we have $v^T A^T A u = v_I^T [A_I^T A_I] u_I$. By RIP, the eigenvalues $\lambda_i = 1 + \mu_i$ of the symmetric matrix $Q = A_I^T A_I$ are in-between $1 - \delta$ and $1 + \delta$; representing u_I and v_I by vectors w and z of their coordinates in the orthonormal eigenbasis of Q , we get $|v^T A^T A u| = |\sum_i \lambda_i w_i z_i| = |\sum_i w_i z_i + \sum_i \mu_i w_i z_i| \leq |w^T z| + \delta \|w\|_2 \|z\|_2$. It remains to note that $w^T z = u_I^T v_I = 0$ and $\|w\|_2 = \|u\|_2$, $\|z\|_2 = \|v\|_2$.

We have

$$\begin{aligned} & \|Ax^1\|_2 \|Ax\|_2 \geq [x^1]^T A^T A x = \|Ax^1\|_2^2 + \sum_{j=2}^q [x^1]^T A^T A x^j \\ & \geq \|Ax^1\|_2^2 - \delta \sum_{j=2}^q \|x^1\|_2 \|x^j\|_2 \quad [\text{by } (*)] \\ & \geq \|Ax^1\|_2^2 - \delta s^{-1/2} \|x^1\|_2 \sum_{j=2}^q \|x^{j-1}\|_1 \geq \|Ax^1\|_2^2 - \delta s^{-1/2} \|x^1\|_2 \|x\|_1 \\ \Rightarrow & \|Ax^1\|_2^2 \leq \|Ax^1\|_2 \|Ax\|_2 + \delta s^{-1/2} \|x^1\|_2 \|x\|_1 \\ \Rightarrow & \|x^1\|_2 = \frac{\|Ax^1\|_2}{\|Ax^1\|_2} \|Ax^1\|_2 \leq \frac{\|Ax^1\|_2}{\|Ax^1\|_2} \|Ax\|_2 + \delta s^{-1/2} \left(\frac{\|x^1\|_2}{\|Ax^1\|_2} \right)^2 \|x\|_1 \\ \Rightarrow & \|x\|_{s,2} = \|x^1\|_2 \leq \frac{1}{\sqrt{1-\delta}} \|Ax\|_2 + \frac{\delta s^{-1/2}}{1-\delta} \|x\|_1 \quad (!) \\ & [\text{by RIP}(\delta, 2s)] \end{aligned}$$

and we see that the pair $\left(H = \frac{s^{-1/2}}{\sqrt{1-\delta}}I_m, \|\cdot\|_2\right)$ satisfies $\mathbf{Q}_2\left(s, \frac{\delta}{1-\delta}\right)$, as claimed in Proposition 1.8.i. Moreover, when $q \geq 2$, $\kappa > 0$, and integer $t \geq 1$ satisfy $t \leq s$ and $\kappa t^{1/q-1} \geq \frac{\delta s^{-1/2}}{1-\delta}$, by (!) we have

$$\|x\|_{t,q} \leq \|x\|_{s,q} \leq \|x\|_{s,2} \leq \frac{1}{\sqrt{1-\delta}}\|Ax\|_2 + \kappa t^{1/q-1}\|x\|_1,$$

or, equivalently,

$$\begin{aligned} 1 \leq t &\leq \min \left[\left[\frac{\kappa(1-\delta)}{\delta} \right]^{\frac{q}{q-1}}, s^{\frac{q-2}{2q-2}} \right] s^{\frac{q}{2q-2}} \\ \Rightarrow (H = \frac{t^{-\frac{1}{q}}}{\sqrt{1-\delta}}I_m, \|\cdot\|_2) &\text{ satisfies } \mathbf{Q}_q(t, \kappa), \end{aligned}$$

as required in Proposition 1.12.i.

Next, we have

$$\begin{aligned} \|x^1\|_1 \|A^T Ax\|_\infty &\geq [x^1]^T A^T Ax = \|Ax^1\|_2^2 + \sum_{j=2}^q [x^1]^T A^T Ax^j \\ &\geq \|Ax^1\|_2^2 - \delta s^{-1/2} \|x^1\|_2 \|x\|_1 \text{ [exactly as above]} \\ \Rightarrow \|Ax^1\|_2^2 &\leq \|x^1\|_1 \|A^T Ax\|_\infty + \delta s^{-1/2} \|x^1\|_2 \|x\|_1 \\ \Rightarrow (1-\delta) \|x^1\|_2^2 &\leq \|x^1\|_1 \|A^T Ax\|_\infty + \delta s^{-1/2} \|x^1\|_2 \|x\|_1 \text{ [by RIP}(\delta, 2s)\text{]} \\ &\leq s^{1/2} \|x^1\|_2 \|A^T Ax\|_\infty + \delta s^{-1/2} \|x^1\|_2 \|x\|_1 \\ \Rightarrow \|x\|_{s,2} = \|x^1\|_2 &\leq \frac{s^{1/2}}{1-\delta} \|A^T Ax\|_\infty + \frac{\delta}{1-\delta} s^{-1/2} \|x\|_1 \end{aligned} \quad (!!)$$

and we see that the pair $\left(H = \frac{1}{1-\delta}A, \|\cdot\|_\infty\right)$ satisfies the condition $\mathbf{Q}_2\left(s, \frac{\delta}{1-\delta}\right)$, as required in Proposition 1.8.ii. Moreover, when $q \geq 2$, $\kappa > 0$, and integer $t \geq 1$ satisfy $t \leq s$ and $\kappa t^{1/q-1} \geq \frac{\delta}{1-\delta} s^{-1/2}$, we have by (!!)

$$\|x\|_{t,q} \leq \|x\|_{s,q} \leq \|x\|_{s,2} \leq \frac{1}{1-\delta} s^{1/2} \|A^T Ax\|_\infty + \kappa t^{1/q-1} \|x\|_1,$$

or, equivalently,

$$\begin{aligned} 1 \leq t &\leq \min \left[\left[\frac{\kappa(1-\delta)}{\delta} \right]^{\frac{q}{q-1}}, s^{\frac{q-2}{2q-2}} \right] s^{\frac{q}{2q-2}} \\ \Rightarrow (H = \frac{s^{\frac{1}{2}t^{-\frac{1}{q}}}}{1-\delta}A, \|\cdot\|_\infty) &\text{ satisfies } \mathbf{Q}_q(t, \kappa), \end{aligned}$$

as required in Proposition 1.12.ii. \square

1.5.5 Proof of Proposition 1.10

(i): Let $\bar{H} \in \mathbf{R}^{m \times N}$ and $\|\cdot\|$ satisfy $\mathbf{Q}_\infty(s, \kappa)$. Then for every $k \leq n$ we have

$$|x_k| \leq \|\bar{H}^T Ax\| + s^{-1} \kappa \|x\|_1,$$

or, which is the same by homogeneity,

$$\min_x \{ \|\bar{H}^T Ax\| - x_k : \|x\|_1 \leq 1 \} \geq -s^{-1} \kappa.$$

In other words, the optimal value Opt_k of the conic optimization problem¹³

$$\text{Opt}_k = \min_{x,t} \{t - [e^k]^T x : \|\bar{H}^T A x\| \leq t, \|x\|_1 \leq 1\},$$

where $e^k \in \mathbf{R}^n$ is k -th basic orth, is $\geq -s^{-1}\kappa$. Since the problem clearly is strictly feasible, this is the same as saying that the dual problem

$$\max_{\mu \in \mathbf{R}, g \in \mathbf{R}^n, \eta \in \mathbf{R}^n} \{-\mu : A^T \bar{H} \eta + g = e^k, \|g\|_\infty \leq \mu, \|\eta\|_* \leq 1\},$$

where $\|\cdot\|_*$ is the norm conjugate to $\|\cdot\|$,

$$\|u\|_* = \max_{\|h\| \leq 1} h^T u,$$

has a feasible solution with the value of the objective $\geq -s^{-1}\kappa$. It follows that there exist $\eta = \eta^k$ and $g = g^k$ such that

$$\begin{aligned} (a) : e^k &= A^T h^k + g^k, \\ (b) : h^k &:= \bar{H} \eta^k, \|\eta^k\|_* \leq 1, \\ (c) : \|g^k\|_\infty &\leq s^{-1}\kappa. \end{aligned} \tag{1.60}$$

Denoting $H = [h^1, \dots, h^n]$, $V = I - H^T A$, we get

$$\text{Col}_k[V^T] = e^k - A^T h^k = g^k,$$

implying that $\|\text{Col}_k[V^T]\|_\infty \leq s^{-1}\kappa$. Since the latter inequality is true for all $k \leq n$, we conclude that

$$\|\text{Col}_k[V]\|_{s,\infty} = \|\text{Col}_k[V^T]\|_\infty \leq s^{-1}\kappa, \quad 1 \leq k \leq n,$$

whence, by Proposition 1.9, $(H, \|\cdot\|_\infty)$ satisfies $\mathbf{Q}_\infty(s, \kappa)$. Moreover, for every $\eta \in \mathbf{R}^m$ and every $k \leq n$ we have, in view of (b) and (c),

$$|[h^k]^T \eta| = |[\eta^k]^T \bar{H}^T \eta| \leq \|\eta^k\|_* \|\bar{H}^T \eta\|,$$

whence $\|H^T \eta\|_\infty \leq \|\bar{H}^T \eta\|$.

Now let us prove the ‘‘In addition’’ part of the proposition. Let $H = [h_1, \dots, h_n]$ be the contrast matrix specified in this part. We have

$$|[I_m - H^T A]_{ij}| = |[e^i]^T - h_i^T A]_j| \leq |[e^i]^T - h_i^T A|_\infty = \|e^i - A^T h_i\|_\infty \leq \text{Opt}_i,$$

implying by Proposition 1.9 that $(H, \|\cdot\|_\infty)$ does satisfy the condition $\mathbf{Q}_\infty(s, \kappa_*)$ with $\kappa_* = s \max_i \text{Opt}_i$. Now assume that there exists a matrix H' which, taken along with some norm $\|\cdot\|$, satisfies the condition $\mathbf{Q}_\infty(s, \kappa)$ with $\kappa < \kappa_*$, and let us lead this assumption to a contradiction. By the already proved first part of Proposition 1.10, our assumption implies that there exists an $m \times n$ matrix $\bar{H} = [\bar{h}_1, \dots, \bar{h}_n]$ such that $\|\text{Col}_j[I_n - \bar{H}^T A]\|_\infty \leq s^{-1}\kappa$ for all $j \leq n$, implying that $\|[e^i]^T - \bar{h}_i^T A]_j| \leq s^{-1}\kappa$ for all i and j , or, which is the same, $\|e^i - A^T \bar{h}_i\|_\infty \leq s^{-1}\kappa$ for all i . Due to the origin of Opt_i , we have $\text{Opt}_i \leq \|e^i - A^T \bar{h}_i\|_\infty$ for all i ,

¹³For a summary on conic programming, see Section 4.1.

and we arrive at $s^{-1}\kappa_* = \max_i \text{Opt}_i \leq s^{-1}\kappa$, that is, $\kappa_* \leq \kappa$, which is a desired contradiction.

It remains to prove (1.33), which is just an exercise on LP duality: denoting by e an n -dimensional all-ones vector, we have

$$\begin{aligned} \text{Opt}_i &:= \min_h \|e^i - A^T h\|_\infty = \min_{h,t} \{t : e^i - A^T h \leq te, A^T h - e^i \leq te\} \\ &= \max_{\lambda, \mu} \{\lambda_i - \mu_i : \lambda, \mu \geq 0, A[\lambda - \mu] = 0, \sum_i \lambda_i + \sum_i \mu_i = 1\} \\ &\quad \text{[LP duality]} \\ &= \max_{x: \lambda - \mu} \{x_i : Ax = 0, \|x\|_1 \leq 1\} \end{aligned}$$

where the concluding equality follows from the fact that vectors x representable as $\lambda - \mu$ with $\lambda, \mu \geq 0$ satisfying $\|\lambda\|_1 + \|\mu\|_1 = 1$ are exactly vectors x with $\|x\|_1 \leq 1$. \square

1.5.6 Proof of Proposition 1.13

Let H satisfy (1.38). Since $\|v\|_{s,1} \leq s^{1-1/q} \|v\|_{s,q}$, it follows that H satisfies for some $\alpha < 1/2$ the condition

$$\|\text{Col}_j[I_n - H^T A]\|_{s,1} \leq \alpha, 1 \leq j \leq n, \quad (1.61)$$

whence, as we know from Proposition 1.9,

$$\|x\|_{s,1} \leq s \|H^T A x\|_\infty + \alpha \|x\|_1 \quad \forall x \in \mathbf{R}^n.$$

It follows that $s \leq m$, since otherwise there exists a nonzero s -sparse vector x with $Ax = 0$; for this x , the inequality above cannot hold true.

Let us set $\bar{n} = 2m$, so that $\bar{n} \leq n$, and let \bar{H} and \bar{A} be the $m \times \bar{n}$ matrices comprised of the first $2m$ columns of H and A , respectively. Relation (1.61) implies that the matrix $V = I_{\bar{n}} - \bar{H}^T \bar{A}$ satisfies

$$\|\text{Col}_j[V]\|_{s,1} \leq \alpha < 1/2, 1 \leq j \leq \bar{n}. \quad (1.62)$$

Now, since the rank of $\bar{H}^T \bar{A}$ is $\leq m$, at least $\bar{n} - m$ singular values of V are ≥ 1 , and therefore the squared Frobenius norm $\|V\|_F^2$ of V is at least $\bar{n} - m$. On the other hand, we can upper-bound this squared norm as follows. Observe that for every \bar{n} -dimensional vector f one has

$$\|f\|_2^2 \leq \max \left[\frac{\bar{n}}{s^2}, 1 \right] \|f\|_{s,1}^2. \quad (1.63)$$

Indeed, by homogeneity it suffices to verify the inequality when $\|f\|_{s,1} = 1$; besides, we can assume w.l.o.g. that the entries in f are nonnegative, and that $f_1 \geq f_2 \geq \dots \geq f_{\bar{n}}$. We have $f_s \leq \|f\|_{s,1}/s = \frac{1}{s}$; in addition, $\sum_{j=s+1}^{\bar{n}} f_j^2 \leq (\bar{n} - s) f_s^2$. Now, due to $\|f\|_{s,1} = 1$, for fixed $f_s \in [0, 1/s]$ we have

$$\sum_{j=1}^s f_j^2 \leq f_s^2 + \max_t \left\{ \sum_{j=1}^{s-1} t_j^2 : t_j \geq f_s, j \leq s-1, \sum_{j=1}^{s-1} t_j = 1 - f_s \right\}.$$

The maximum on the right-hand side is the maximum of a convex function

over a bounded polytope; it is achieved at an extreme point, that is, at a point where one of the t_j is equal to $1 - (s - 1)f_s$, and all remaining t_j are equal to f_s . As a result,

$$\sum_j f_j^2 \leq [(1 - (s - 1)f_s)^2 + (s - 1)f_s^2] + (\bar{n} - s)f_s^2 \leq (1 - (s - 1)f_s)^2 + (\bar{n} - 1)f_s^2.$$

The right-hand side in the latter inequality is convex in f_s and thus achieves its maximum over the range $[0, 1/s]$ of allowed values of f_s at an endpoint, yielding $\sum_j f_j^2 \leq \max[1, \bar{n}/s^2]$, as claimed.

Applying (1.63) to the columns of V and recalling that $\bar{n} = 2m$, we get

$$\|V\|_F^2 = \sum_{j=1}^{2m} \|\text{Col}_j[V]\|_2^2 \leq \max\left[1, \frac{2m}{s^2}\right] \sum_{j=1}^{2m} \|\text{Col}_j[V]\|_{s,1}^2 \leq 2\alpha^2 m \max\left[1, \frac{2m}{s^2}\right].$$

The left hand side in this inequality, as we remember, is $\geq \bar{n} - m = m$, and we arrive at

$$m \leq 2\alpha^2 m \max[1, 2m/s^2].$$

Since $\alpha < 1/2$, this inequality implies $2m/s^2 \geq 2$, whence $s \leq \sqrt{m}$.

It remains to prove that when $m \leq n/2$, the condition $\mathbf{Q}_\infty(s, \kappa)$ with $\kappa < 1/2$ can be satisfied only when $s \leq \sqrt{m}$. This is immediate: by Proposition 1.10, assuming $\mathbf{Q}_\infty(s, \kappa)$ satisfiable, there exists an $m \times n$ contrast matrix H such that $|[I_n - H^T A]_{ij}| \leq \kappa/s$ for all i, j , which, by the already proved part of Proposition 1.13, is impossible when $s > \sqrt{m}$. \square

Index

- $O(1)$, xvi
 Diag, xv
 Erfc, 16
 ErfcInv, 17
 Risk, 41
 $\mathbf{E}_\xi\{ \cdot \}$, $\mathbf{E}_{\xi \sim P}\{ \cdot \}$, $\mathbf{E}\{ \cdot \}$, xvi
 $\mathbf{Q}_q(s, \kappa)$ -condition, 12
 links with RIP, 24
 tractability when $q = \infty$, 22
 verifiable sufficient conditions
 for, 20
 $\mathbf{R}^n, \mathbf{R}^{m \times n}$, xv
 $\mathbf{R}_+, \mathbf{R}_+^n$, xvi
 \mathbf{S}^n , xv
 \mathbf{S}_+^n , xvi
 \mathcal{A}^* , xv
 $\mathcal{N}(\mu, \Theta)$, xvi
 $\mathcal{R}_k[\cdot], \mathcal{R}_k^*[\cdot], \mathcal{S}_\ell[\cdot], \mathcal{S}_\ell^*[\cdot], \dots$, 275
 dg, xv
 ℓ_1 minimization, *see* Compressed Sensing
 Sensing
 Poisson(\cdot), xvi
 Uniform(\cdot), xvi
 $\int_\Omega f(\xi) \Pi(d\xi)$, xvi
 $\lambda[\cdot]$, 276
 $\succeq, \succ, \preceq, \prec$, xvi
 $\xi \sim P, \xi \sim p(\cdot)$, xvi
 s -goodness, 9
 $\| \cdot \|_p$, xvi
 $\| \cdot \|_{2,2}$, xvi

 Bisection estimate, 198
 near-optimality of, 202

 closeness relation, 59
 Compressed Sensing, 3–6
 via ℓ_1 minimization, 6–17
 imperfect, 11
 validity of, 8
 verifiable sufficient validity conditions, 17–26
 verifiable sufficient validity conditions, limits of performance, 24
 via penalized ℓ_1 recovery, 14
 via regular ℓ_1 recovery, 13
 conditional quantile, 198
 cone
 dual, 262
 Lorentz, 262
 regular, 262
 semidefinite, 262
 conic
 problem, 263
 dual of, 263
 programming, 262, 265
 Conic Duality Theorem, 264
 conic hull, 263
 contrast matrix, *see* nullspace property quantification
 Cramer-Rao risk bound, 347–350, 354–355, 570

 detector, 65
 affine, 123
 in simple observation schemes, 83
 quadratic, 139
 risks of, 65
 structural properties, 65

 ellitope, 265–266
 calculus of, 299–302, 429
 estimation
 of N -convex functions, 193–211
 of linear form, 185, 211–222
 from repeated observations, 215–217
 of sub-Gaussianity parameters, 217–222
 of sub-Gaussianity parameters, direct product case, 219

- of quadratic form, 222–232
 - Gaussian case, 222–228
 - Gaussian case, consistency, 227
 - Gaussian case, construction, 224
 - sub-Gaussian case, 228, 232
 - sub-Gaussian case, construction, 228
- family of distributions
 - regular/simple, 124–132
 - calculus of, 126
 - examples of, 125
 - spherical, 53
 - cap of, 53
- function
 - N -convex, 197
 - examples of, 197
- Gaussian mixtures, 53
- generalized linear model, 415
- GLM, *see* generalized linear model
- Hellinger affinity, 83
- Hypothesis Testing
 - change detection
 - via quadratic lifting, 149–156
 - of multiple hypotheses, 58–64
 - in simple observation schemes, 87–105
 - up to closeness, 59, 91
 - via Euclidean separation, 62–64
 - via repeated observations, 95
 - of unions, 87
 - problem's setting, 41
 - sequential, 105–113
 - test, 42
 - detector-based, 65
 - detector-based, limits of performance, 70
 - detector-based, via repeated observations, 66
 - deterministic, 42
 - partial risks of, 45
 - randomized, 42
 - simple, 42
 - total risk of, 45
 - two-point lower risk bound, 46
 - via affine detectors, 132–139
 - via Euclidean separation, 49–58
 - and repeated observations, 55
 - majority test, 56
 - multiple hypotheses case, 62–64
 - pairwise, 50
 - via quadratic lifting, 139
 - Gaussian case, 139–145
 - sub-Gaussian case, 145–149
 - via repeated observations, 42
- inequality
 - Cramer-Rao, 349, 572
- lemma
 - on Schur Complement, *see* Schur Complement Lemma
- LMI, xvi
- logistic regression, 414–415
- matrices
 - notation, xv
 - sensing, 1
- MD, *see* Measurement Design
- Measurement Design, 113–123
 - simple case
 - discrete o.s., 118
 - Gaussian o.s., 122
 - Poisson o.s., 121
- Mutual Incoherence, 23
- norm
 - conjugate, 279
 - Shatten, 305
 - Wasserstein, 340, 559
- Nullspace property, 9, 10
 - quantification of, 11
- o.s., *see* observation scheme
- observation scheme
 - discrete, 77, 85
 - Gaussian, 74, 84
 - Poisson, 75, 84
 - simple, 72–87
 - K -th power of, 85
 - definition of, 73
 - direct product of, 77
- PET, *see* Positron Emission Tomography

- Poisson Imaging, 75
- polyhedral estimate, 385–414
- Positron Emission Tomography, 75

- Rademacher random vector, 363
- regular data, 124
- repeated observations
 - quasi-stationary, 44
 - semi-stationary, 43
 - stationary, 42
- Restricted Isometry Property, 20
- RIP, *see* Restricted Isometry Property
- risk
 - $\text{Risk}(\mathcal{T}|H_1, \dots, H_L)$, 45
 - $\text{RiskOpt}_{\Pi, \|\cdot\|}[\mathcal{X}]$, 290
 - $\text{Risk}[\hat{x}(\cdot)|\mathcal{X}]$, 260
 - Risk_ϵ^* , 234
 - $\text{Risk}_\epsilon^{\mathcal{C}}(\mathcal{T}|H_1, \dots, H_L)$, 60
 - $\text{Risk}_\epsilon^{\mathcal{C}}(\mathcal{T}|H_1, \dots, H_L)$, 60
 - $\text{Risk}_\epsilon^{\text{opt}}(K)$, 220
 - $\text{Risk}_\ell(\mathcal{T}|H_1, \dots, H_L)$, 45
 - $\text{Risk}_\epsilon(\hat{g}(\cdot)|G, \mathcal{X}, v, \mathcal{A}, \mathcal{H}, \mathcal{M}, \Phi)$, 212
 - $\text{Risk}_\pm[\phi|\mathcal{P}]$, $\text{Risk}[\phi|\mathcal{P}_1, \mathcal{P}_2]$, 65
 - $\text{Risk}_{\text{tot}}(\mathcal{T}|H_1, \dots, H_L)$, 45
 - $\text{Risk}_{\mathcal{H}}[\hat{x}_*|\mathcal{X}]$, 298
 - $\text{Risk}_{\text{opt}}[\mathcal{X}]$, 260
 - $\text{Risk}_{\mathcal{H}, \|\cdot\|}[\hat{x}|\mathcal{X}]$, 295
 - \mathcal{C} -, 60
 - \mathcal{H} -, 295
 - ϵ -, 212
 - $\text{Risk}_{\Pi, \|\cdot\|}[\hat{x}|\mathcal{X}]$, 279
 - $\text{Risk}_{\Pi, \mathcal{H}, \|\cdot\|}[\hat{x}|\mathcal{X}]$, 299
 - in Hypothesis Testing
 - partial, 45
 - total, 45
 - up to closeness, 59
 - minimax, 266
 - ϵ -, 220
 - of detector, 65
 - of simple test, 45
- SA, *see* Stochastic Approximation
- SAA, *see* Sample Average Approximation
- saddle point
 - convex-concave saddle point problem, 79
- Sample Average Approximation, 419–421
- Schur Complement Lemma, 265
- semidefinite relaxation
 - on ellitope
 - tightness of, 274
 - on spectratope
 - tightness of, 277
- signal estimation, *see* signal recovery
- signal recovery
 - linear, 267
 - on ellitope, 267–271
 - on ellitope, near-optimality of, 271–274
 - on spectratope, 277–291
 - on spectratope under uncertain-but-bounded noise, 291–297
 - on spectratope under uncertain-but-bounded noise, near-optimality of, 297
 - on spectratope, near-optimality of, 277, 290
 - problem setting, 1, 260
- sparsity, s -sparsity, 3
- spectratope, 275
 - calculus of, 299–302, 429
 - examples of, 276
- Stochastic Approximation, 421–424

- test, *see* Hypothesis Testing test
- theorem
 - Sion-Kakutani, 81

- vectors
 - notation, xv