

Compressive Sensing

Collection Editors:

Mark A. Davenport

Richard Baraniuk

Ronald DeVore

Compressive Sensing

Collection Editors:

Mark A. Davenport
Richard Baraniuk
Ronald DeVore

Authors:

Wai Lam Chan
Mark A. Davenport
Ronald DeVore
Marco F. Duarte
Jason Laska

Michael A Lexa
Chris Rozell
Shriram Sarvotham
Michael Wakin

Online:

< <http://cnx.org/content/col10458/1.1/> >

C O N N E X I O N S

Rice University, Houston, Texas

This selection and arrangement of content as a collection is copyrighted by Mark A. Davenport, Richard Baraniuk, Ronald DeVore. It is licensed under the Creative Commons Attribution 2.0 license (<http://creativecommons.org/licenses/by/2.0/>).
Collection structure revised: September 21, 2007
PDF generated: October 26, 2012
For copyright and attribution information for the modules contained in this collection, see p. 35.

Table of Contents

Introduction	1
1 Encoding of Analog Signals	
1.1 The Shannon-Whitaker Sampling Theorem	3
1.2 Optimal Encoding	5
1.3 Kolmogorov Entropy	6
1.4 Optimal Encoding of Bandlimited Signals	7
1.5 Stable Signal Representations	9
2 Sparse Signals and Sparse Approximation	
2.1 Preliminaries	13
2.2 New Signal Models	13
2.3 Sparse Approximation and ℓ_p Spaces	15
2.4 Thresholding and Greedy Bases	17
2.5 Greedy Algorithms	19
3 Compressive Sensing	
3.1 Compressive Sensing	23
3.2 Gelfand n-widths	25
3.3 Instance Optimality	26
3.4 The Restricted Isometry Property	27
3.5 The Nullspace Property	28
3.6 Optimality and the MRIP	29
3.7 Summary	32
Attributions	35

Introduction¹

Throughout this course, we shall be interested in the analog to digital conversion of signals $f(t)$, $t \in \mathbb{R}$. We shall always assume $f \in L_2$ and usually assume additional properties of f in order to get meaningful results. In particular, we want to study two mappings: the encoding of f into bit streams and the decoding of the bit streams into approximations or estimates of f ,

$$\begin{aligned} E : f &\rightarrow \text{bits streams} && \text{(Encoder)} \\ D : \text{bits streams} &\rightarrow \bar{f} && \text{(Decoder)} \end{aligned} \tag{1}$$

where \bar{f} is the approximation of f defined by $\bar{f} := D(E(f))$. In general, $\bar{f} \neq f$, so we shall need some way of quantifying how well \bar{f} approximates f . Normally, the distortion between is measured by some norm $\|f - \bar{f}\|$. Typical choices include:

$$\begin{aligned} \text{the } L_2 \text{ norm} \quad \|f\|_{L_2} &:= \left(\int |f(t)|^2 dt \right)^{1/2} \\ \text{the } L_\infty \text{ norm} \quad \|f\|_{L_\infty} &:= \sup_t |f(t)| \\ \text{the } L_p \text{ norm} \quad \|f\|_{L_p} &:= \left(\int |f(t)|^p dt \right)^{1/p} \end{aligned} \tag{2}$$

¹This content is available online at <<http://cnx.org/content/m15136/1.1/>>.

Chapter 1

Encoding of Analog Signals

1.1 The Shannon-Whitaker Sampling Theorem¹

The classical theory behind the encoding analog signals into bit streams and decoding bit streams back into signals, rests on a famous sampling theorem which is typically referred to as the Shannon-Whitaker Sampling Theorem. In this course, this sampling theory will serve as a benchmark to which we shall compare the new theory of compressed sensing.

To introduce the Shannon-Whitaker theory, we first define the class of bandlimited signals. A bandlimited signal is a signal whose Fourier transform only has finite support. We shall denote this class as B_A and define it in the following way:

$$B_A := \{f \in L_2(\mathbb{R}) : \hat{f}(\omega) = 0, |\omega| \geq A\pi\}. \quad (1.1)$$

Here, the Fourier transform of f is defined by

$$\hat{f}(\omega) := \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} f(t) e^{-i\omega t} dt. \quad (1.2)$$

This formula holds for any $f \in L_1$ and extends easily to $f \in L_2$ via limits. The inversion of the Fourier transform is given by

$$f(t) := \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} \hat{f}(\omega) e^{i\omega t} d\omega. \quad (1.3)$$

Theorem 1.1: Shannon-Whitaker Sampling Theorem

If $f \in B_A$, then f can be uniquely determined by the uniformly spaced samples $f\left(\frac{n}{A}\right)$ and in fact, is given by

$$f(t) = \sum_{n \in \mathbb{Z}} f\left(\frac{n}{A}\right) \operatorname{sinc}(\pi(A t - n)), \quad (1.4)$$

where $\operatorname{sinc}(t) = \frac{\sin t}{t}$.

Proof:

It is enough to consider $A = 1$, since all other cases can be reduced to this through a simple change of variables. Because $f \in B_{A=1}$, the Fourier inversion formula takes the form

$$f(t) = \frac{1}{\sqrt{2\pi}} \int_{-\pi}^{\pi} \hat{f}(\omega) e^{i\omega t} d\omega. \quad (1.5)$$

¹This content is available online at <<http://cnx.org/content/m15146/1.2/>>.

Define $F(\omega)$ as the 2π periodization of \hat{f} ,

$$F(\omega) := \sum_{n \in \mathbb{Z}} \hat{f}(\omega - 2n\pi). \quad (1.6)$$

Because $F(\omega)$ is periodic, it admits a Fourier series representation

$$F(\omega) = \sum_{n \in \mathbb{Z}} c_n e^{-in\omega}, \quad (1.7)$$

where the Fourier coefficients c_n given by

$$\begin{aligned} c_n &= \frac{1}{2\pi} \int_{-\pi}^{\pi} F(\omega) e^{in\omega} d\omega \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \hat{f}(\omega) e^{in\omega} d\omega. \end{aligned} \quad (1.8)$$

By comparing ((1.8)) with ((1.5)), we conclude that

$$c_n = \frac{1}{\sqrt{2\pi}} f(n). \quad (1.9)$$

Therefore by plugging ((1.9)) back into ((1.6)), we have that

$$F(\omega) = \frac{1}{\sqrt{2\pi}} \sum_{n \in \mathbb{Z}} f(n) e^{-in\omega}. \quad (1.10)$$

Now, because

$$\hat{f}(\omega) = F(\omega) \chi_{[-\pi, \pi]} = \frac{1}{\sqrt{2\pi}} \sum_{n \in \mathbb{Z}} f(n) e^{-in\omega} \chi_{[-\pi, \pi]}, \quad (1.11)$$

and because of the facts that

$$\begin{aligned} \mathcal{F}(\chi_{[-\pi, \pi]}) &= \frac{1}{\sqrt{2\pi}} \text{sinc}(\pi\omega) \quad \text{and} \\ \mathcal{F}(g(t-n)) &= e^{-in\omega} \mathcal{F}(g(t)), \end{aligned} \quad (1.12)$$

we conclude

$$f(t) = \sum_{n \in \mathbb{Z}} f(n) \text{sinc}(\pi(t-n)). \quad (1.13)$$

Comments:

1. (Good news) The set $\{\text{sinc}(\pi(t-n))\}_{n \in \mathbb{Z}}$ is an orthogonal system and therefore, has the property that the L_2 norm of the function and its Fourier coefficients are related by,

$$\|f\|_{L_2}^2 = 2\pi \sum_{n \in \mathbb{Z}} |f(n)|^2 \quad (1.14)$$

2. (Bad news) The representation of f in terms of sinc functions is not a stable representation, i.e.

$$\sum_{n \in \mathbb{Z}} |\text{sinc}(\pi(t-n))| \approx \sum_{n \in \mathbb{Z}} \frac{1}{|t-n|+1} \rightarrow \text{divergences} \quad (1.15)$$

1.2 Optimal Encoding²

We shall consider now the encoding of signals on $[-T, T]$ where $T > 0$ is fixed. Ultimately we shall be interested in encoding classes of bandlimited signals like the class B_A . However, we begin the story by considering the more general setting of encoding the elements of any given compact subset K of a normed linear space X . One can determine the best encoding of K by what is known as the Kolmogorov entropy of K in X .

To begin, let us consider an encoder-decoder pair (E, D) . E maps K to a finite stream of bits. D maps a stream of bits to a signal in X . This is illustrated in Figure 1.1. Note that many functions can be mapped onto the same bitstream.




Figure 1.1: Illustration of encoding and decoding.

Define the distortion d for this encoder-decoder by

$$d(K, E, D, X) := \sup_{f \in K} \|f - D(Ef)\|_{\underline{X}}. \quad (1.16)$$

Let $n(K, E) = \sup_{f \in K} \#Ef$ where $\#Ef$ is the number of bits in the bitstream Ef . Thus n is the maximum length of the bitstreams for the various $f \in K$. There are two ways we can define optimal encoding:

1. Prescribe ε , the maximum distortion that we are willing to tolerate. For this ε , find the smallest $n_\varepsilon(K, X) := \inf_{(E, D)} \{n(K, E) : d(K, E, D, X) \leq \varepsilon\}$. This is the smallest bit budget under which we could encode all elements of K to distortion ε .
2. Prescribe N : find the smallest distortion $d(K, E, D, X)$ over all E, D with $n(K, E) \leq N$. This is the best encoding performance possible with a prescribed bit budget.

There is a simple mathematical solution to these two encoding problems based on the notion of Kolmogorov Entropy.

²This content is available online at <http://cnx.org/content/m15139/1.1/>.

1.3 Kolmogorov Entropy³

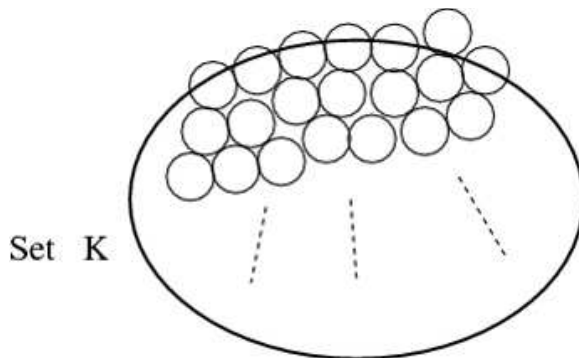


Figure 1.2: Coverings of K by balls of radius ε .

Given $\varepsilon > 0$, and the compact set K , consider all coverings of K by balls of radius ε , as shown in Figure 1.2. In other words,

$$K \subseteq \bigcup_{i=1}^N b(f_i, \varepsilon). \quad (1.17)$$

Let $N_\varepsilon := \inf \{N : \text{over all such covers}\}$. $N_\varepsilon(K)$ is called the covering number of K . Since it depends on X and K , we write it as $N_\varepsilon = N_\varepsilon(K, X)$.

Rule 1.1: Kolmogorov entropy

The Kolmogorov entropy, denoted by $H_\varepsilon(K, X)$, of the compact set K in X is defined as the logarithm of the covering number:

$$H_\varepsilon(K, X) = \log N_\varepsilon(K, X). \quad (1.18)$$

The Kolmogorov entropy solves our problem of optimal encoding in the sense of the following theorem.

Theorem 1.2:

For any compact set $K \subset X$, we have $n_\varepsilon(K, X) = \lceil H_\varepsilon(K, X) \rceil$, where $\lceil \cdot \rceil$ is the ceiling function.

Proof:

Sketch: We can define an encoder-decoder as follows To encode: Say $f \in K$. Just specify which ball it is covered by. Because the number of balls is $N_\varepsilon(K, \overline{X})$, we need at most $\lceil \log N_\varepsilon(K, \overline{X}) \rceil$ bits to specify any such ball ball.

To decode: Just take the center of the ball specified by the bitstream.

It is now easy to see that this encoder-decoder pair is optimal in either of the senses given above.

□

The above encoder is not practical. However, the Kolmogorov entropy tells us the best performance we can expect from any encoder-decoder pair. Kolmogorov entropy is defined in the deterministic setting. It is the analogue of the Shannon entropy which is defined in a stochastic setting.

³This content is available online at <http://cnx.org/content/m15137/1.1/>.

1.4 Optimal Encoding of Bandlimited Signals⁴

We now turn back to the encoding of signals. We are interested in encoding the set

$$B_A(M) = \{f \in B_A : |f(t)| \leq M, t \in \mathbb{R}\} \quad (1.19)$$

where M is arbitrary but fixed. We shall restrict our discussion to the case where distortion is measured in $L_\infty[-T, T]$ where $T > 0$ is arbitrary but fixed. Then, $B_A(M)$ is a compact subset of L_∞ : $B_A(M) \subseteq L_\infty[-T, T]$.




Figure 1.3: Sample points $\frac{n}{\lambda A}$ are chosen in the interval $[-T(1 + \delta), T(1 + \delta)]$.

We shall sketch how one can construct an asymptotically optimal encoder/decoder for B_A . The details for this construction can be found in .

We know $\hat{f}(\omega) = 0$ for $|\omega| \geq A\pi$, and $|f| \leq M$. How can we encode f in practice? We begin by choosing $\lambda = \lambda(T) > 1$ (see Figure 1.3) which will represent a slight oversampling factor we shall utilize. Given a target distortion $\varepsilon > 0$, we choose k so that $2^{-k-1} < \varepsilon \leq 2^{-k}$. Given f , we shall encode f by first taking samples $f\left(\frac{n}{\lambda A}\right)$ for $\frac{n}{\lambda A} \in [-T(1 + \delta), T(1 + \delta)]$ where $\delta(T) > 0$. In other words, we sample f on a slightly larger interval than $[-T, T]$. For each sample $f\left(\frac{n}{\lambda A}\right)$, we shall use the first $k + k_0(T)$ bits of its binary expansion. In other words, our encoder takes f and the samples $f\left(\frac{n}{\lambda A}\right)$ and then assigns to $f\left(\frac{n}{\lambda A}\right)$ the first $k + k_0(T)$ bits of this number.

To decode, the receiver would take the bits and construct the approximation $\bar{f}\left(\frac{n}{\lambda A}\right)$ to $f\left(\frac{n}{\lambda A}\right)$ from the bits provided. Notice that we have the accuracy

$$\left|f\left(\frac{n}{\lambda A}\right) - \bar{f}\left(\frac{n}{\lambda A}\right)\right| \leq 2^{-k-k_0} \cdot M. \quad (1.20)$$

We utilize the function g_λ satisfying () to define

$$\bar{f}(t) = \sum_{n \in N_T} \bar{f}\left(\frac{n}{\lambda A}\right) g_\lambda(\lambda A t - n), \quad (1.21)$$

where

$$N_T := \{n : -T(1 + \delta) \leq \frac{n}{\lambda A} \leq T(1 + \delta)\}. \quad (1.22)$$

We then have

$$\begin{aligned} |f(t) - \bar{f}(t)| &\leq \sum_{n \in N_T} \left|f\left(\frac{n}{\lambda A}\right) - \bar{f}\left(\frac{n}{\lambda A}\right)\right| \cdot |g_\lambda(\lambda A t - n)| \\ &\quad + \sum_{\left|\frac{n}{\lambda A}\right| > T(1 + \delta)} \left|f\left(\frac{n}{\lambda A}\right)\right| \cdot |g_\lambda(\lambda A t - n)| \end{aligned} \quad (1.23)$$

⁴This content is available online at <<http://cnx.org/content/m15140/1.2/>>.

The term $|f(\frac{n}{\lambda A}) - \bar{f}(\frac{n}{\lambda A})|$ that appears in the first summation in ((1.23)) is bounded by $M \cdot 2^{-k-k_0}$. The term $|f(\frac{n}{\lambda A})|$ that appears in the second summation in the same equation is bounded by M . Therefore,

$$|f(t) - \bar{f}(t)| \leq \sum_{n \in N_T} M \cdot 2^{-k-k_0} \cdot |g_\lambda(\lambda At - n)| + \sum_{|\frac{n}{\lambda A}| > T(1+\delta)} M \cdot |g_\lambda(\lambda At - n)| =: S_1 + S_2 \quad (1.24)$$

We can estimate S_1 by

$$\begin{aligned} S_1 &= \sum_{n \in N_T} M \cdot 2^{-k-k_0} \cdot |g_\lambda(\lambda At - n)| \\ &\leq M \cdot 2^{-k-k_0} \cdot \sum_n |g_\lambda(\lambda At - n)| \\ &\leq M \cdot C_0(\lambda) \cdot 2^{-k-k_0} \quad (\text{because } g(\cdot) \text{ decays fast}) \end{aligned} \quad (1.25)$$

Therefore, if we choose k_0 sufficiently large, then $S_1 \leq M \cdot C_0(\lambda) \cdot 2^{-k-k_0} \leq \frac{\varepsilon}{2}$. The second summation S_2 can also be bounded by $\varepsilon/2$ by using the fast decay of the function g_λ (see ()).

To make the encoder/decoder specific we need to precisely define δ and λ . It turns out that the best choices (in terms of bit rate performance on the class B_A) depend on T . But $\delta_T \rightarrow 0$ and $\lambda_T \rightarrow 1$ as $T \rightarrow \infty$. Recall that Shannon sampling requires $2T\lambda A$ samples. Since our encoder/decoder uses $k + k_0$ bits per sample, the total number of bits is $(k + k_0) \cdot 2\lambda AT(1 + \delta)$, and so coding will require roughly k bits per Shannon sample.

This encoder/decoder can be proven to be optimal in the sense of averaged performance as we shall now describe. The average of performance of optimal encoding is defined by

$$\lim_{T \rightarrow \infty} \frac{n_\varepsilon(B_A(M), L_\infty[-T, T])}{2T} \quad (1.26)$$

If we replace the optimal bit rate n_ε in ((1.26)) by the number of bits required by our encoder/decoder then the resulting limit will be the same as that in ((1.26)).

In summary, to encode band limited signals on an interval $[-T, T]$, an optimal strategy is to sample at a slightly higher rate than Nyquist and on a slightly large interval than $[-T, T]$. Each sample should then be quantized by using the binary expansion of the sample. In this way, for an investment of k bits per Nyquist rate sample, we get a distortion of 2^{-k} .

To get a feel for the number of bits required by such an encoder, let us say $A = 10^6$ (signals band limited to 1Mhz). Say $T = 24$ hours $\approx 10^5$ seconds, and $k = 10$ bits. Then, $A \cdot k \cdot 2T = 10^6 \cdot 10 \cdot 10^5 = 10^{12}$ bits. This is too BIG!

The above encoding is known as Pulse Coded Modulation (PCM). In practice, people frequently use another encoder called Sigma-Delta Modulation. Instead of oversampling just slightly, Sigma Delta oversamples a lot and then assign only one (or a few) bits per sample.

Why is Sigma-Delta preferred to PCM in practice? There are two reasons commonly given:

1. Getting accurate samples, quantization, etc. is not practical because of noise. For better accuracy, we need more expensive hardware.
2. Noise shaping. In Sigma-Delta, the distortion is higher but the distortion is spread over frequencies outside of the desired range.

In PCM, the distortion decays exponentially (like 2^{-k}), whereas for Sigma-Delta, the distortion decays like a polynomial (like $\frac{1}{k^m}$). Although the distortion decays faster in PCM, the distortion in Sigma-Delta is spread outside the desired frequency range.

1.5 Stable Signal Representations⁵

To fix the instability of the Shannon representation, we assume that the signal is slightly more bandlimited than before

$$\hat{f}(\omega) = 0 \quad \text{for } |\omega| \geq \pi - \delta, \quad \delta > 0, \quad (1.27)$$

and instead of using $\chi_{[-\pi, \pi]}$, we multiply by another function $\hat{g}(\omega)$ which is very similar in form to the characteristic function, but decays at its boundaries in a smoother fashion (i.e. it has more derivatives). A candidate function \hat{g} is sketched in Figure 1.4.




Figure 1.4: Sketch of \hat{g} .

Now, it is a property of the Fourier transform that an increased smoothness in one domain translates into a faster decay in the other. Thus, we can fix our instability problem, by choosing \hat{g} so that \hat{g} is smooth and $\hat{g}(\omega) = 1$, $|\omega| \leq \pi - \delta$ and $\hat{g} = 0$, $|\omega| > \pi$. By choosing the smoothness of g suitably large, we can, for any given $m \geq 1$, choose g to satisfy

$$|g(t)| \leq \frac{C}{(|t| + 1)^m} \quad (1.28)$$

for some constant $C > 0$.

Using such a \hat{g} , we can rewrite () as

$$\hat{f}(\omega) = F(\omega) \hat{g}(\omega) = \frac{1}{\sqrt{2\pi}} \sum_{n \in \mathbb{Z}} f(n) e^{-in\omega} \hat{g}(\omega). \quad (1.29)$$

Thus, we have the new representation

$$f(t) = \sum_{n \in \mathbb{Z}} f(n) g(t - n), \quad (1.30)$$

where we gain stability from our additional assumption that the signal is bandlimited on $[-\pi - \delta, \pi - \delta]$.

⁵This content is available online at <<http://cnx.org/content/m15144/1.1/>>.

Does this assumption really hurt? No, not really because if our signal is really bandlimited to $[-\pi, \pi]$ and not $[-\pi - \delta, \pi - \delta]$, we can always take a slightly larger bandwidth, say $[-\lambda\pi, \lambda\pi]$ where λ is a little larger than one, and carry out the same analysis as above. Doing so, would only mean slightly oversampling the signal (small cost).

Recall that in the end we want to convert analog signals into bit streams. Thus far, we have the two representations

$$\begin{aligned} f(t) &= \sum_{n \in \mathbb{Z}} f(n) \operatorname{sinc}(\pi(t - n)), \\ f(t) &= \sum_{n \in \mathbb{Z}} f\left(\frac{n}{\lambda}\right) g(\lambda t - n). \end{aligned} \quad (1.31)$$

Shannon's Theorem tells us that if $f \in B_A$, we should sample f at the Nyquist rate A (which is twice the support of \hat{f}) and then take the binary representation of the samples. Our more stable representation says to slightly oversample f and then convert to a binary representation. Both representations offer perfect reconstruction, although in the more stable representation, one is straddled with the additional task of choosing an appropriate λ .

In practical situations, we shall be interested in approximating f on an interval $[-T, T]$ for some $T > 0$ and not for all time. Questions we still want to answer include

1. How many bits do we need to represent f in $B_{A=1}$ on some interval $[-T, T]$ in the norm $L_\infty[-T, T]$?
2. Using this methodology, what is the optimal way of encoding?
3. How is the optimal encoding implemented?

Towards this end, we define

$$B_A := \{f \in L_2(\mathbb{R}) : \hat{f}(\omega) = 0, |\omega| \geq A\pi\}. \quad (1.32)$$

Then for any $f \in B_A$, we can write

$$f = \sum_n f\left(\frac{n}{A}\right) \cdot \operatorname{sinc} \pi(At - n). \quad (1.33)$$




Figure 1.5: Fourier transform of $g_\lambda(\cdot)$.

In other words, samples at $0, \pm \frac{1}{A}, \pm \frac{2}{A}, \dots$ are sufficient to reconstruct f . Recall also that $\operatorname{sinc}(x) = \frac{\sin(x)}{x}$ decays poorly (leading to numerical instability). We can overcome this problem by slight over-sampling. Say we over-sample by a factor $\lambda > 1$. Then, we can write

$$f = \sum f\left(\frac{n}{\lambda A}\right) g_\lambda(\lambda At - n). \quad (1.34)$$

Hence we need samples at $0, \pm \frac{1}{\lambda A}, \pm \frac{2}{\lambda A}, \text{ etc.}$ What is the advantage? Sampling more often than necessary buys us stability because we now have a choice for $g_\lambda(\cdot)$. If we choose $g_\lambda(\cdot)$ infinitely differentiable whose Fourier transform looks as shown in Figure 1.5 we can obtain

$$|g_\lambda(t)| \leq \frac{c_{\lambda,k}}{(1+|t|)^k}, \quad k = 1, 2, \dots \quad (1.35)$$

and therefore $g_\lambda(\cdot)$ decays very fast. In other words, a sample's influence is felt only locally. Note however, that over-sampling generates basis functions that are redundant (linearly dependent), unlike the integer translates of the sinc (\cdot) function.




Figure 1.6: To reconstruct signals in $[-T, T]$, the sampling interval is $[-cT, cT]$.

If we restrict our reconstruction to t in the interval $[-T, T]$, we will only need samples only from $[-cT, cT]$, for $c > 1$ (see Figure 1.6), because the distant samples will have little effect on the reconstruction in $[-T, T]$.

Chapter 2

Sparse Signals and Sparse Approximation

2.1 Preliminaries¹

We previously described Shannon's Theorem plus encoding: the Nyquist sampling rate is the minimal required sampling rate to recover the entire class of bandlimited signals. We have seen that this sampling rate may be prohibitively large for broadband signals. We see a way to improve upon this situation: we will pose a different model for the signals which is more restrictive than the assumption that the signals are bandlimited. Fortunately, there are several real world scenarios in which one knows much more information about the signals of interest. For example, they may be written in terms of very few fundamental building blocks (such as sine waves or chirps). This leads us to define new signal classes based on notions of sparsity and seek to determine if we can improve on sampling and encoding in this new setting.

Let us define the general setting for this section. Let \mathbf{X} be a Banach space of functions. The typical examples are $\mathbf{X} = L_p(\mathbb{R}), L_p(\mathbb{R}^d), L_p(-T, T)$, $1 \leq p \leq \infty$. We denote the norm on X by $\|\cdot\|_{\mathbf{X}}$. We define a dictionary \mathcal{D} as any collection of functions $\mathcal{D} \subseteq \mathbf{X}$ such that $\|g\|_{\mathbf{X}} = 1$ for all $g \in \mathcal{D}$, i.e. all the elements of the dictionary are normalized. While the definition is very broad, in practice dictionaries usually have more structure. Some examples include $\mathcal{D} = B$, a basis for \mathbf{X} , such as (i) the Fourier basis on $[-\pi, \pi]$, (ii) a wavelet basis,² (iii) redundant families of waveforms of the form $\psi_{a,b,\sigma} = e^{-a(t-b)^2} e^{i\sigma x}$, i.e. $\mathcal{D} = \{\psi_{a,b,\sigma}\}_{a,b,\sigma}$, and (iv) wavelet packets.

Rule 2.1:

We define the class of n -sparse signals as $\Sigma_n := \Sigma_n(\mathcal{D}) = \{s = \sum_{g \in \Lambda} c_g g, \Lambda \subseteq \mathcal{D}, \#\Lambda \leq n\}$. We also say that s has sparsity n in \mathcal{D} if $s \in \Sigma_n(\mathcal{D})$, i.e. if it can be written as the linear combination of n functions from \mathcal{D} . We note that Σ_n is not a linear space; we instead have $\Sigma_n + \Sigma_n \subseteq \Sigma_{2n}$.

2.2 New Signal Models³

We now wish to consider new model classes for signals. Towards this end, let $\{\psi_j\}_{j=1}^{\infty}$ be an orthonormal basis for $L_2(-T, T)$. Thus for $f \in L_2$ we can write $f = \sum_{j=1}^{\infty} c_j(f) \psi_j$ where $(c_j(f)) \in \ell_2$. We will now build an encoder and decoder and analyze its performance on compact sets K . For example, we might want to encode signals in the space

$$X_p = \{f : (c_j(f)) \in \ell_p\}, 0 \leq p \leq 2$$

¹This content is available online at <<http://cnx.org/content/m15142/1.1/>>.

²Wavelet basis form orthonormal systems for $L_2(I)$.

³This content is available online at <<http://cnx.org/content/m15138/1.1/>>.

with norm

$$\|f\|_{X_p} := \|(c_j(f))\|_{\ell_p}.$$

However, in this space the unit ball, $U(X_p)$ is not compact. To get a compact set we need more structure on the sequence (c_j) . Hence we define

$$Y^\alpha := \{f : |c_n(f)| \leq n^{-\alpha}, n = 1, 2, \dots\}$$

and we define the norm in this space as $\|f\|_{Y^\alpha} :=$ the smallest c such that this holds. We now take

$$K = U(X_p) \cap U(Y^\alpha)$$

to get a compact set. Notice that when $\alpha > 0$ is small the requirement for membership in Y^α is very mild.

Next, suppose that we choose a target distortion level $\epsilon = 2^{-m}$. Given f , let

$$\Lambda_k := \Lambda_k(f) = \{j \in \{0, \dots, N\} : 2^{-k-1} \leq |c_j(f)| < 2^{-k}\}$$

for $0 \leq k \leq M$, where $M := \lceil \frac{2m}{2-p} \rceil$. We then choose N as the smallest integer so that

$$N^{-\alpha} \leq 2^{-M}$$

and thus

$$\log N \leq Cm.$$

It follows from the requirement that $f \in Y^\alpha$ that $\Lambda_k \subset \{1, \dots, N\}$ for each $0 \leq k \leq M$.

Recall that

$$\#\Lambda_k 2^{-(k-1)p} \leq \sum_{c_j \in \Lambda_k} |c_j|^p \leq \|f\|_{X_p}^p.$$

Since $f \in U(X_p) \cap U(Y^\alpha)$,

$$\#\Lambda_k \leq \|f\|_{\ell_p} 2^{(k+1)p} \leq 2^{(k+1)p}.$$

Hence, the total number of indices in all of the Λ_k , $0 \leq k \leq M$, is $O(2^{Mp})$.

To encode, for each f , we can send the following bits:

- Send $\log n$ bits to identify each index in Λ_k , for $0 \leq k \leq M$. This will require a total of $O(\log N 2^{Mp})$ bits.
- Send one bit to identify the sign of $c_j(f)$ for each $j \in \Lambda_k$, $0 \leq k \leq M$. This will require $O(2^{Mp})$ bits.
- Send m bits to describe each $c_j(f)$, $j \in \Lambda_k$, for $0 \leq k \leq M$. This will require $O(m 2^{Mp})$ bits.

Thus the total number of bits used in the encoding is $O(m 2^{Mp})$.

Notice that for each $j \in \Lambda_k$, $0 \leq k \leq M$, we can recover each $c_j(f)$ by

$$\bar{c}_j = \pm \sum_{i=0}^m b_i 2^{-k-i}$$

where the sign is given by the sign bit. It follows that $|c_j(f) - \bar{c}_j| \leq 2^{-m-k}$ for every such coefficient. Here we have used the fact that knowing that $j \in \Lambda_k$ means that the first nonzero binary bit of $c_j(f)$ is the k -th bit.

To decode we simply set

$$\bar{f} = \sum_{k=0}^M \sum_{j \in \Lambda_k} \bar{c}_j \psi_j$$

We now analyze the error we have incurred in such an encoding. The square of the error will consist of two parts. The first corresponds to the $j \in \Lambda_k$, $0 \leq k \leq M$. For each such j we have $|c_j(f) - \bar{c}_j| \leq 2^{-m-k}$ and so the total square error for this is

$$\leq C \sum_{k=1}^M 2^{kp} 2^{-2m} 2^{-2k} \leq c 2^{-2m}$$

because $p \leq 2$. The second part of the error corresponds to all the coefficients which have magnitude $\leq 2^{-M}$. We have that this sum does not exceed $\sum_{|c_j| > 2^{-M}} |c_j|^2 \leq 2^{-M(2-p)} \sum_{j=1}^{\infty} |c_j|^p \leq 2^{-2m}$. Thus the total error we incur is $O(2^{-m})$.

In summary, by allocating $O\left(m 2^{\frac{m}{1/p-1/2}}\right)$ bits we achieve distortion $C 2^{-m}$. Equivalently, by allocating $n \log n$ bits, we achieve distortion $C n^{-(1/p-1/2)}$.

Rule 2.1:

This is within a logarithmic factor of the optimal encoding given by Kolmogorov entropy of the class $U(X_p) \cap Y^\alpha$. A slightly more careful argument can remove this logarithm.

Example 2.1: The Wavelet Basis

In the method above we failed to achieve the optimal performance because of the cost involved in identifying which indices were in each Λ_k . We will now describe a method that can do better, using the Haar basis for $L_2[0, 1]$. Thus, we first define the scaling function

$$\phi := \chi_{[0,1]}.$$

Next, we define the mother wavelet

$$\psi := \chi_{[0, \frac{1}{2}]} - \chi_{[\frac{1}{2}, 1]}.$$

We then define the remaining wavelets recursively. They are obtained by dilations and shifts of the mother wavelet on dyadic intervals:

$$\psi_J := \frac{2^k}{2} \psi_{[0,1]}(2^k x - j)$$

where $J = [j 2^{-k}, (j+1) 2^{-k}]$ are dyadic intervals. We denote by D_+ the collection of all dyadic intervals contained in $[0, 1]$. Then, the collection of functions $\{\phi\} \cup \{\psi_J\}_{J \in D_+}$ forms an orthonormal basis for $L_2[0, 1]$.

A key property of wavelets is that a tree structure can be placed on the coefficients due to the use of dyadic intervals in their construction. Thus, let

$$T_k := \{j : |c_j| \geq 2^{-k}\}$$

and

$$T_{k+1} - T_k = \Lambda_k.$$

We define \bar{T}_k as the smallest tree containing T_k . Given any binary tree of size n , we can encode the tree with at most $O(n)$ bits, in the process outperforming the encoder described in above.

2.3 Sparse Approximation and ℓ_p Spaces⁴

We now look at how well $f \in \mathbf{X}$ can be approximated by n functions in the dictionary \mathcal{D} .

⁴This content is available online at <http://cnx.org/content/m15143/1.1/>.

Rule 2.2:

We define the error of n -term approximation of f by the elements of the dictionary \mathcal{D} as

$$(1) \quad \sigma_n(f)_{\mathbf{X}} := \sigma_n(f, \mathcal{D})_{\mathbf{X}} := \inf_{s \in \Sigma_n} \|f - s\|_{\mathbf{X}}.$$

We also define the class of r -smooth signals in \mathcal{D} as

$$(2) \quad \mathcal{A}^r := \mathcal{A}^r(\mathcal{D}) := \{f \in X, \sigma_n(f) \leq Mn^{-r} \text{ for some } M\}$$

with the corresponding norm $\|f\|_{\mathcal{A}^r} = \sup_{n=1,2,\dots} n^r \sigma_n(f)_{\mathbf{X}}$.

In general, the larger r is, the 'smoother' the function $s \in \mathcal{A}^r(\mathcal{D})_{\mathbf{X}}$. Note also that $\mathcal{A}^r \subseteq \mathcal{A}^{r'}$ if $r > r'$. Given f , let $r(f) = \sup\{r : f \in \mathcal{A}^r\}$ be a measure of the "smoothness" of f , i.e. a quantification of compressibility.

Let $\mathbf{X} = H$, a Hilbert space⁵ such as $\mathbf{X} = L_2(\mathbb{R})$, and assume $\mathcal{D} = B$ - an orthonormal basis on \mathbf{X} ; i.e. if $B = \{\varphi_i\}_i$, then $\langle \varphi_i, \varphi_j \rangle = \delta_{i,j}$, where $\delta_{i,j}$ is the Kronecker delta. This also means that each $f \in \mathbf{X}$ has an expansion $f = \sum_j c_j(f) \varphi_j$, where $c_j(f) = \langle f, \varphi_j \rangle$. We also have $\|f\|_{\mathbf{X}}^2 = \sum_{j=1}^{\infty} |c_j(f)|^2$.

Recall the definition of ℓ_p spaces: let $(a_j) \in \mathbb{R}$; then $(a_j) \in \ell_p$ if $\|(a_j)\|_{\ell_p} < \infty$ with $\|(a_j)\|_{\ell_p} = \left(\sum_j |a_j|^p\right)^{1/p}$ for $p < \infty$ and $\|(a_j)\|_{\ell_p} = \sup_j |a_j|$ for $p = \infty$. We also recall that for L_p spaces on compact sets, $L_p \subset L_{p'}$ if $p > p'$. The opposite is true for ℓ_p spaces: $\ell_p \subset \ell_{p'}$ if $p < p'$. Hence, the smaller the value of p is, the "smaller" ℓ_p is.

Example 2.2

Does there exist a sequence (a_j) with $\|(a_j)\|_{\ell_1} = \sum_j |a_j| < \infty$ but with $\|(a_j)\|_{\ell_p} = \left(\sum_j |a_j|^p\right)^{1/p} = \infty$ for all $0 < p < 1$? Consider the sequence $a_n = \frac{1}{n(\log n)^{1+\delta}}$. We see that $(a_n) \in \ell_1$ but $\|(a_n)\|_{\ell_p} = \infty$ for all $0 < p < 1$.

Lemma 2.1:

A sequence (a_n) is in ℓ_p if the sorted magnitudes of the a_n decay faster than $n^{-\frac{1}{p}}$.

Define a_n^* as the element of the sequence (a_n) with the n^{th} largest magnitude, and denote (a_n^*) as the decreasing rearrangement of (a_n) . It is easy to show that $k(a_k^*)^p \leq \sum_n (a_n)^p$ for all k ; also, if $(a_n) \in \ell_p$, then $a_k^* \leq \|(a_n)\|_{\ell_p} k^{-\frac{1}{p}}$.

Rule 2.3:

A sequence (a_n) is in weak ℓ_p , denoted $(a_n) \in w\ell_p$, if $a_k^* \leq Mk^{-\frac{1}{p}}$. We also define the quasinorm⁶ $\|(a_n)\|_{w\ell_p}$ as the smallest $M > 0$ such that $a_k^* \leq Mk^{-\frac{1}{p}}$ for each k .

Example 2.3

The sequence $a_n = \frac{1}{n}$ is in weak ℓ_1 but not in ℓ_1 .

Lemma 2.2:

For p, p' such that $p' > p$, we have $\ell_p \subset w\ell_p \subset \ell_{p'}$.

Theorem 2.1:

Let $\mathcal{D} = B$ be an orthonormal basis for the Hilbert space $\mathbf{X} = H$. For $f \in \mathbf{X}$ with representation in $B = [\varphi_1, \varphi_2, \dots]$ as $f = \sum_n c_n(f) \varphi_n$, we have $f \in \mathcal{A}^r(B)_{\mathbf{X}}$ if and only if the sequence $(c_n(f)) \in w\ell_r$, with $\frac{1}{r} = r + \frac{1}{2}$. Moreover, there exist $C_0, C'_0 \in \mathbb{R}$ such that $C'_0 \|(c_n(f))\|_{w\ell_r} \leq \|f\|_{\mathcal{A}^r} \leq C_0 \|(c_n(f))\|_{w\ell_r}$.

Example

Let $r = \frac{1}{2}$. $f \in \mathcal{A}^{\frac{1}{2}}$ if and only if $(c_n(f)) \in w\ell_r$, i.e. if $c_n^*(f) \leq Mn^{-1} = \frac{M}{n}$.

⁵A Hilbert space is a complete inner product space with the norm induced by the inner product

⁶A quasinorm is has the properties of a norm except that the triangle inequality is replaced by the condition $\|x + y\| \leq C_0(\|x\| + \|y\|)$ for some absolute constant C_0 .

Proof:

We prove the converse statement; the forward statement proof is left to the reader. We would like to show that if $(c_n(f)) \in w\ell_\tau$, then $f \in \mathcal{A}^r$, with $r = \frac{1}{\tau} - \frac{1}{2}$. The best n -term approximation of f in B is of the form $s = \sum_{k \in \Lambda} a_k \varphi_k$, $\#\Lambda \leq n$. Therefore, we have:

$$\begin{aligned}
 \sigma_n(f)_{\mathbf{X}} &= \inf_{s \in \Sigma_n} \|f - s\|_{\mathbf{X}} = \inf_{s \in \Sigma_n} \left\| \sum_{k \in \Lambda} (c_k(f) - a_k) \varphi_k + \sum_{k \notin \Lambda} c_k(f) \varphi_k \right\|_{\mathbf{X}} \\
 (3) \quad &= \inf_{s \in \Lambda} \sum_{k \in \Lambda} (c_k(f) - a_k)^2 + \sum_{k \notin \Lambda} (c_k(f))^2 = \sum_{k=n+1}^{\infty} |c_k^*(f)|^2 \\
 &\leq M^2 \sum_{k=n+1}^{\infty} k^{-\frac{2}{\tau}} \leq M^2 \sum_{k=n+1}^{\infty} k^{-2r-1} \text{ (since } (c_n(f)) \in w\ell_p),
 \end{aligned}$$

where $M := \| (c_n(f)) \|_{w\ell_p}$.

We prove the converse statement; the forward statement proof is left to the reader. We would like to show that if $(c_n(f)) \in w\ell_\tau$, then $f \in \mathcal{A}^r$, with $r = \frac{1}{\tau} - \frac{1}{2}$. The best n -term approximation of f in B is of the form $s = \sum_{k \in \Lambda} a_k \varphi_k$, $\#\Lambda \leq n$. Therefore, we have:

$$\begin{aligned}
 \sigma_n(f)_{\mathbf{X}} &= \inf_{s \in \Sigma_n} \|f - s\|_{\mathbf{X}} = \inf_{s \in \Sigma_n} \left\| \sum_{k \in \Lambda} (c_k(f) - a_k) \varphi_k + \sum_{k \notin \Lambda} c_k(f) \varphi_k \right\|_{\mathbf{X}} \\
 (3) \quad &= \inf_{s \in \Lambda} \sum_{k \in \Lambda} (c_k(f) - a_k)^2 + \sum_{k \notin \Lambda} (c_k(f))^2 = \sum_{k=n+1}^{\infty} |c_k^*(f)|^2 \\
 &\leq M^2 \sum_{k=n+1}^{\infty} k^{-\frac{2}{\tau}} \leq M^2 \sum_{k=n+1}^{\infty} k^{-2r-1} \text{ (since } (c_n(f)) \in w\ell_p),
 \end{aligned}$$

where we define $C = \frac{\lambda_0}{r}$. Using this result in the earlier statement, we get

$$(5) \quad \sum_{k=n+1}^{\infty} |c_k^*(f)|^2 \leq CM^2 n^{-2r} \leq M^2 z n^{-r};$$

this implies by definition that $(c_k(f)) \in \mathcal{A}^r$.

2.4 Thresholding and Greedy Bases⁷

We shall next discuss some notions related to best n -term approximation.

2.4.1 Thresholding

- Let \mathbb{X} be a Hilbert space. Given f , let $\Lambda_\varepsilon(f) = \{j : |c_j(f)| > \varepsilon\}$. The thresholding operator T_ε is defined by

$$T_\varepsilon f = \sum_{j \in \Lambda_\varepsilon(f)} C_j(f) \phi_j. \quad (2.1)$$

It is easy to see that for each ε , $T_\varepsilon f$ is the best approximation to f using N terms where N is the cardinality of Λ_ε :

$$\|f - T_\varepsilon f\|_{\mathbb{X}} = \sigma_N(f)_{\mathbb{X}}. \quad (2.2)$$

Thresholding is easily implemented on a computer.

- The thresholding scheme above can be generalized if \mathbb{X} is not a Hilbert space provided the dictionary has some specific structure. For example, when
 - The dictionary is the wavelet basis and $\mathbb{X} = L_p$, $1 < p < \infty$.
 - $\mathbb{X} = l_p$ and the dictionary is the canonical basis $\delta_j = \phi_j$. ex: $(0, 0, \dots, 1, 0)$
 - For a general Banach space \mathbb{X} and the dictionary (ϕ_j) is a greedy basis.

⁷This content is available online at <http://cnx.org/content/m15147/1.1/>.

2.4.2 Greedy Bases

We briefly describe the notion of greedy basis.

Rule 2.4:

Given \mathbb{X} , we say (ϕ_j) is a greedy basis for \mathbb{X} if for each $\varepsilon > 0$,

$$\|f - T_\varepsilon f\|_{\mathbb{X}} \leq C(\mathbb{X}) \sigma_N(f)_{\mathbb{X}} \quad (2.3)$$

where N is the cardinality of Λ_ε .

Rule 2.5:

A basis ϕ_j is said to be unconditional if

$$\left\| \sum \pm c_j \phi_j \right\|_{\mathbb{X}} \leq C \left\| \sum c_j \phi_j \right\|_{\mathbb{X}} \quad (2.4)$$

or equivalently

$$\left\| \sum c_j \phi_j \right\|_{\mathbb{X}} \leq C \left\| \sum d_j \phi_j \right\|_{\mathbb{X}} \quad \text{where} \quad |c_j| \leq |d_j|. \quad (2.5)$$

This is an older concept from functional analysis. In words, this definition says that if the terms c_j are rearranged, the series $\sum c_j \phi_j$ will still converge. This is not generally true for all bases.

Rule 2.6:

A basis ϕ_j is said to be democratic if

$$\left\| \sum_{j \in \Lambda} \phi_j \right\| \leq C \left\| \sum_{j \in \Lambda'} \phi_j \right\|, \quad (2.6)$$

where the cardinality of Λ' equals the cardinality of Λ .

Rule 2.2:

(ϕ_j) greedy \leftrightarrow (ϕ_j) is both unconditional and democratic.

Some examples involving the last two definitions:

- The fourier basis in L_p is not democratic, but is unconditional for $1 < p < \infty$.
- The wavelet basis contains both of these properties, and is therefore greedy.

Theorem 2.2:

If $\mathbb{X} = L_p$ has (ϕ_j) greedy, $B = \{\phi_j\}$, $f = \sum_{j=1}^{\infty} c_j(f) \phi_j$, $c_j(f) = \langle f, \psi_j \rangle$ where ψ_j is a dual basis,

$$f \in \mathcal{A}^r(B) \leftrightarrow (c_j(f)) \in w_{l_\tau}, \quad \frac{1}{\tau} = r + \frac{1}{p} \quad (2.7)$$

and

$$\|f\|_{\mathcal{A}^r} \approx \|c_j(f)\|_{l_\tau} \quad (2.8)$$

Let us now consider a specific setting that we shall be concerned with a lot in this course. We shall examine some of the concepts we have introduced in the finite dimensional space of all sequence (points) in \mathbb{R}^N . Recall that we can put many different norms on this space including the ℓ_p norms and the weak ℓ_p norms.

Rule 2.3:

Given a vector $x = (x_1, x_2, \dots, x_N) \in \mathbb{R}^N$. The best approximation to x from Σ_n in the ℓ_p norm is to take the vector in Σ_n which shares the n largest values of x . Its error of approximation satisfies

$$\sigma_n(x)_{\ell_p} \leq Cn^{-r} \|x\|_{w_{l_\tau}}, \quad \frac{1}{\tau} = r + \frac{1}{p} \quad (2.9)$$

Rule 2.4:

$$\|x\|_{w_{l_\tau}} \leq \|x\|_{l_\tau}, \quad \frac{1}{\tau} = r + \frac{1}{p}.$$

Rule 2.1:

For $p = 1$ and $r = 3$, $\sigma_n(x)_{\ell_1} \leq Cn^{-3} \|x\|_{w_{l_\tau}}$ and $\frac{1}{\tau} = 4$. In words, this equation shows what kind of τ is needed for a given decay rate (or given some τ , what kind of decay rate will be achieved) to approximate with certain ability.

Rule 2.2:

Show $\sigma_n(x)_{l_p} \leq Cn^{-r} \|x\|_{l_\tau}$ holds with $C = 1$.

Proof: Let $\Lambda_n := \{i : |x_i| \text{ largest}\}$,

$$\sigma_n(x)_{l_p}^p = \sum_{i \notin \Lambda_n} |x_i|^p \quad (2.10)$$

$$\leq \sum_{i \notin \Lambda_n} |x_i|^{p-\tau} |x_i|^\tau \quad (2.11)$$

$$\leq \left(\|x\|_{w_{l_\tau}} n^{-\frac{1}{\tau}} \right)^{p-\tau} \left(\sum |x_i|^\tau \right) \quad (2.12)$$

$$\leq \|x\|_{l_\tau}^{p-\tau} \|x\|_{l_\tau}^\tau n = n^{-r p} \|x\|_{l_\tau}^p \quad (2.13)$$

and so

$$\sigma_n(x)_{l_p} \leq n^{-r p} \|x\|_{l_\tau}^p \quad (2.14)$$

$$\sigma_n(x)_{l_p} \leq n^{-r} \|x\|_{l_\tau}$$

Rule 2.5:

For $\mathbb{X} = L_p$, $\{\phi_j\}$ a wavelet basis, we can say wavelet coefficients of f are in l_p is equivalent to f is in a certain Besov class (roughly speaking f has r derivatives and $f^{(r)} \in L_\tau$). We refer the reader to for precise formulations of results of this type.

2.5 Greedy Algorithms⁸

We now turn to the questions of generating good approximations for n term approximation from a general dictionary. We shall assume that the dictionary \mathcal{D} is complete in the Hilbert space \mathbb{H} . This means that every element in \mathbb{H} can be approximated arbitrarily well by linear combinations of the elements of \mathcal{D} . Since the dictionary is no longer an orthogonal basis as was considered above, we need to revisit how to find good n term approximations. Because of redundancy within the dictionary, we cannot simply pick the largest coefficients as we saw with a basis. Greedy algorithms are a method to generate good n term approximations.

⁸This content is available online at <http://cnx.org/content/m15134/1.1/>.

1. General Greedy Algorithm Given f , we want to generate an n -term approximation to f .

$$f \rightarrow s = \sum_{s=1}^n c_j g_j \quad (2.15)$$

The general steps are as follows:

- Initialize: (approximation) $s_0 = 0$, (residual) $r_0 = f$, approximation collection $\Lambda_0 = \emptyset$
- Search \mathcal{D} for some $g \in \mathcal{D}$, then add g to the set Λ .
- Use $\{g_1, g_2, \dots, g_n\}$ to find new approximation for s_n .

At stage n , we have s_n , $r_n = f - s_n$, and $\Lambda_n = \{g_1, g_2, \dots, g_n\}$. There are many types of greedy algorithms. We describe the three most common in the case \mathbb{S} is a Hilbert space. However, there are analogues of these for L_p .

2. Pure Greedy Algorithm (PGA) Note: >From r_n choose $g_{n+1} := \operatorname{argmax} | \langle r_n(f), g \rangle |$ (the g that causes the largest inner product).

$$s_{n+1} = s_n + \langle r_n(f), g \rangle g \quad (2.16)$$

$$r_{n+1} = f - s_{n+1} = f - s_n - \langle f - s_n, g \rangle g = f - s_{n+1} \quad (2.17)$$

This method is similar to a steepest decent algorithm for decreasing the error.

3. Orthogonal Greedy Algorithm (OGA) >From r_n choose $g_{n+1} := \operatorname{argmax} | \langle r_n(f), g \rangle |$ as in the PGA.

$$V_{n+1} := \operatorname{sp}\{g_1, g_2, \dots, g_{n+1}\} \quad (2.18)$$

$$s_{n+1} := P_{V_{n+1}} f = \sum_{j=1}^{n+1} \alpha_j g_j \quad (2.19)$$

where P_V denotes the orthogonal projection onto the space V . We can find $s_{n+1} = P_{V_{n+1}} f$ by solving the linear system of equations

$$\left\langle \sum_{j=1}^{n+1} \alpha_j g_j, g_k \right\rangle = \langle f, g_k \rangle. \quad (2.20)$$

Then, $r_{n+1} = f - s_{n+1}$.

4. Relaxed Greedy Algorithm (RGA) >From r_n choose g_{n+1} in some way (for example, our earlier methods) and then define

$$s_{n+1}(f) = \alpha s_n + \beta g_{n+1} \quad (2.21)$$

Unlike PGA, here we do not make a full step in the correct direction. For example, one way to proceed is to define

$$\operatorname{arginf}_{\alpha, \beta, g} \| f - \alpha s_n + \beta g_n \| =: \alpha^*, \beta^*, g^* \quad (2.22)$$

This type of greedy algorithm is known to perform the best as compared with the previous two.

2.5.1 Measuring Performance

Given \mathbb{X} , \mathcal{D} , it is not practical to minimize $\sigma_n(f)_{\mathbb{X}}$ by searching over all the possibilities. The greedy approximation gives an n -term solution with less computation, but does it perform well?

Let

$$\mathcal{L}^1(\mathcal{D}) := \{f \in \mathbb{X} : \sum_{g \in \mathcal{D}} c_g g, \sum_{g \in \mathcal{D}} |c_g| \leq M\} \quad (2.23)$$

where the smallest M is the \mathcal{L}^1 norm of f .

Theorem 2.3:

For OGA or RGA as described above, we have

$$\|f - s_n f\| \leq C n^{-\frac{1}{2}} |f|_{\mathcal{L}^1}. \quad (2.24)$$

Rule 2.6:

Remark 5 This is similar to $\sigma_n(x)_{l_2} \leq n^{-\frac{1}{2}} \|x\|_{l_1}$ (n -term approximation) but its not always quite as good.

Chapter 3

Compressive Sensing

3.1 Compressive Sensing¹

We now consider a different setting. Suppose $x \in \mathbb{R}^N$ and we wish to sample x , where taking a sample means the application of a linear functional $\lambda \in \mathbb{R}^N$ to x . Next, we prescribe a budget of n samples, and consider all linear encoders using n samples. We can write these n linear functionals as an $n \times N$ matrix $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}^n$. We then consider a decoder $\Delta : \mathbb{R}^n \rightarrow \mathbb{R}^N$. Our approximation to x is thus $\Delta(\Phi(x))$.

To make the problem precise, we first pick a measure for distortion:

$$\text{error} = \|\Phi(x) - \Delta(\Phi(x))\|_{\ell_p}.$$

We next must make some assumption about x . For example, we can assume that

$$x \in \Sigma_k = \{x : x_i = 0 \text{ for } i \notin \Lambda, \#\Lambda \leq k\},$$

or

$$x \in \ell_\tau$$

or

$$x \in w_{\ell_\tau}$$

We recall our basic problem: a signal $x \in \mathbb{R}^N$ will be “sampled” or “sensed” by applying the n linear projections represented by the columns of the sampling matrix $\Phi_{n \times N}$. The resulting measurements are given in the vector $y \in \mathbb{R}^n$, where $y = \Phi x$. We will assume that $n < N$, meaning that in addition to thinking of the sampling operation Φ as an encoder, we can also view it as a projection to a lower dimensional linear subspace. In either case, we would like the measurements y to preserve as much information about the signal x as possible. To proceed in finding optimal solutions, we must formalize this problem and define how we will measure this information loss.

Moving forward, a critical quantity for us will be the null space of the sampling matrix, $N = N(\Phi) = \{x : \Phi x = 0\}$. Because we are trying to take as few measurements as possible, we will assume that we are taking measurements efficiently so that the rows of Φ are all linearly independent. Therefore, assume that $\text{rank}(\Phi) = n$, implying that N has dimension $(N - n)$. The non-trivial null space of Φ means that it is not an invertible mapping. For any measurement vector y we can define the class of all observable signals that would result in the same measurement, $\mathcal{F}(y) := \{x : \Phi x = y\} \subseteq \mathbb{R}^N$. The class $\mathcal{F}(y)$ can always be written as a sum of a vector in the class and a vector in the null space of the sampling matrix, $\mathcal{F}(y) = x_0 + N$, where $x_0 \in \mathcal{F}(y)$. If we have two vectors $x_0, x_1 \in \mathcal{F}(y)$, then by linearity we know that $\Phi(x_1 - x_0) = 0$. This fact implies that $(x_1 - x_0) \in N$, and consequently that $x_1 \in x_0 + N$.

¹This content is available online at <http://cnx.org/content/m15132/1.1/>.

Associated to our encoder Φ , we shall have to describe a decoder Δ , which is a (not necessarily linear) map $\Delta : \mathbb{R}^n \rightarrow \mathbb{R}^N$. This decoder will take the measurements y and try to recover x as closely as possible, $\bar{x} = \Delta(\Phi x)$. In order to design the best decoder Δ , we must specify our metric for measuring the quality of the estimate \bar{x} . For the moment we shall always think of taking an optimal decoder for the problem at hand. Later we shall discuss specific and concrete decoders.

To begin our discussion of the efficiency of the encoder Φ we consider the following problem.

NOTE: We will fix n and N and try to find Φ and a decoder Δ such that for all input signals in a sparsity class $x \in \Sigma_k$ we can get perfect reconstruction, $\Delta(\Phi x) = x$. We will be interested in determining what the largest value of k is that we can find such an encoder/decoder pair.

An important role in this problem and later problems of compressed sensing is played by certain submatrices of Φ . Given a set $T \subseteq \{1, 2, \dots, N\}$, representing a collection of column indices we define the matrix Φ_T as the one formed from Φ by using the columns from the set T . The matrix Φ_T is a $(n \times \#(T))$ matrix. But sometimes we will also use the same notation Φ_T to denote the matrix obtained from Φ by setting all entries not in the columns of T to zero.

We can now state the following theorem.

Theorem 3.1:

The following statements are all equivalent for any given Φ :

1. There exists a decoder Δ such that $\Delta(\Phi x) = x$ for all $x \in \Sigma_k$.
2. $N(\Phi) \cap \Sigma_{2k} = \{0\}$.
3. Φ_T has rank $\#(T)$ for all T with $\#(T) = 2k$.
4. $\Phi_T^t \Phi_T$ is non-singular (i.e., invertible) for all T with $\#(T) = 2k$.

Proof:

The equivalence of $b \leftrightarrow c \leftrightarrow d$ is simple linear algebra. First let us prove that $a \rightarrow b$. Assume a . Suppose that there was a vector $\eta \in N \cap \Sigma_{2k}$. We know that we can write $\eta = x_0 - x_1$, where x_0, x_1 both have support less than k . We could write η simply as a composite vector with support $2k$ where the first half of η is x_0 and the second half of η is x_1 , $\eta = (x_0 \mid x_1)$. We know that $\eta \in N$, which implies that $\Phi \eta = 0$. It then follows that $\Phi x_0 = \Phi x_1$, and we have as a consequence that $\Delta \Phi x_0 = \Delta \Phi x_1$ and finally that $x_0 = x_1$. This proves that $\eta = 0$ as desired.

Now let us prove that $b \rightarrow a$. Suppose that we have a measurement $\Phi x = y \in \mathbb{R}^n$, where $x \in \Sigma_k$. We will define the decoder $\Delta(y)$ to be the signal $\Psi \in \mathcal{F}(y)$ with smallest support. Since x has support k so will Ψ . We claim that there is no other $x' \in \mathcal{F}(y) \cap \Sigma_k$. Indeed, if x' existed, then $x - x' \in N \cap \Sigma_{2k}$. But, the only vector in $N \cap \Sigma_{2k}$ is zero, implying that $x = x'$. This finally gives us that $\Delta(\Phi x) = x$. \square

Using the previous theorem, we can turn to the question of finding good encoder/decoder pairs.² Given a fixed N , how large can k be, and what is the best Φ ? Given a fixed n , the largest k is $k = \lfloor \frac{n}{2} \rfloor$. Alternatively, we can say that given a fixed k , we need at least $n = 2k$ measurements. Another way to say this is that there exist encoding matrices $\Phi_{2k \times N}$ such that any selection of $2k$ columns are linearly independent. Examples are the DFT matrix or the Vandermonde matrix corresponding to interpolation at distinct points z_1, \dots, z_N .

²A question about whether we will ever really find natural signals in Σ_k brings to mind a story... Once upon a time, a man was floating over the countryside in a hot air balloon. The man in the balloon yelled down to a stranger on the ground and asked "Where am I?" The man on the ground thought for about 5 minutes and then answered "You're in a hot air balloon." The man in the balloon responded with "You must be a mathematician," to which the man on the ground answered "Yes, how did you know?" "Because," replied the man in the balloon, "you had to think a long time before you answered, your answer was very precise, and your answer was completely useless!" So, yes, we may be dealing with a limited model, but we have to crawl before we can walk.

3.2 Gelfand n-widths³

Continuing from last time, we have a signal $x \in \mathbb{R}^N$, an $n \times N$ measurement matrix Φ , and $y = \Phi x \in \mathbb{R}^n$ is the information we draw from x .

We consider decoders Δ mapping $\mathbb{R}^n \rightarrow \mathbb{R}^N$. We have been discussing whether there exists a decoder with certain properties. So for this discussion (about information preservation), we can just think about optimal decoding.

While the previous result on sparse signal recovery is interesting, it is not very robust. What happens if our signal does not have support k ? Can we still obtain meaningful results? The answer is yes. To do so we introduce more general input signal classes $K \subseteq X$ that allows fully supported signals. For example, we will consider the signal class defined by the unit ball

$$K = U(\ell_p^N) = \{x : \|x\|_X \leq 1\}. \quad (3.1)$$

Given an encoder/decoder pair (Φ, Δ) , the worst case error on a set K for that pair will be given by

$$E(K, \Phi, \Delta)_X = \sup_{x \in K} \|x - \Delta(\Phi x)\|_X. \quad (3.2)$$

Finally, using min-max principles we will define the minimum error over all encoder/decoder pairs for a signal class and for a fixed number of measurements n to be

$$E_n(K)_X = \inf_{(\Phi, \Delta): \Phi \text{ is } n \times N} E(K, \Phi, \Delta)_X. \quad (3.3)$$

This measure $E_n(K)_X$ is the best we could do while measuring distortion on the topology of X , using n linear measurements, and using arbitrary decoding.

We will see that these questions are actually equivalent to a classical “ n -width” problem. n -widths have seen a great deal of work over the years by a variety of mathematicians: Kolmogorov, Tikhomirov, Kashin, Gluskin, etc. There are many different flavors of n -widths, but we will study the Gelfand n -width (the least intuitive of the widths).

Rule 3.1: Gelfand n -width

Let $K \subseteq X$ be compact. Given n , the Gelfand width (also called the dual width) is given by

$$d^n(K)_X := \inf_{Y: \text{codim}(Y)=n} \sup\{\|x\|_X : x \in K \cap Y\}. \quad (3.4)$$

where by codimension $(Y) = n$ we mean that Y has dimension $\dim(X) - n$.

In other words, we are looking for the subspace Y that slices through the set K so that the norms of the projected signals are as small as possible. We can now state the following theorem about n -widths:

Theorem 3.2:

Provided that K has the properties (1) $K = -K$ and (2) $K + K = CK$, then

$$d^n(K)_X \leq E_n(K)_X \leq Cd^n(K)_X \quad (3.5)$$

where C is the same constant listed in property (2).⁴

Proof:

We start with the left-hand inequality. We want to take any encoder/decoder pair and use that to construct a Y . So let Φ, Δ be an encoder/decoder. Then simply let $Y = \mathcal{N}(\Phi)$. Now consider

³This content is available online at <<http://cnx.org/content/m15133/1.1/>>.

⁴Clarifying notation: $CK = \{Cx : x \in K\}$ and $K + K = \{x_1 + x_2 : x_1, x_2 \in K\}$.

an $\eta \in K \cap Y$ and note that $\Phi(\eta) = 0$ since $\eta \in Y$. Let $z = \Delta(0)$ be the decoding of 0 (practically speaking, z should be zero itself, but we avoid that assumption in this proof). Then

$$\begin{aligned} \|\eta\|_X &\leq \max(\|\eta - z\|_X, \|\eta + z\|_X) \\ &= \max(\|\eta - \Delta\Phi(\eta)\|_X, \|\eta - \Delta\Phi(\eta)\|_X) \\ &\leq \sup_{x \in K} \|x - \Delta\Phi(x)\|_X \end{aligned} \quad (3.6)$$

where we first employ the triangle inequality, then the fact that multiplying by -1 does not change the norm, then the fact that $K = -K$. So then

$$\sup_{\eta \in K \cap Y} \|\eta\|_X \leq \sup_{x \in K} \|x - \Delta\Phi(x)\|_X. \quad (3.7)$$

Taking the infimum over all Φ, Δ , it follows that

$$d^n(K)_X \leq E_n(K)_X. \quad (3.8)$$

Since Φ is $n \times N$, then $\dim(\mathcal{N}(\Phi)) \geq N - n$.

Now we prove the right-hand inequality. Assume we have a good Y . Suppose Y has codimension $N - n$. Then Y^\perp (the orthogonal complement of Y in \mathbb{R}^N) has dimension n . Let $v_1, v_2, \dots, v_n \in \mathbb{R}^N$ be a basis for Y^\perp . Let Φ be the $n \times N$ matrix obtained by stacking the rows v_1, v_2, \dots, v_n . Then $\mathcal{N}(\Phi) = Y$. Define $\Delta(y) =$ any element of $K \cap \mathcal{F}(y)$ if there is one (otherwise let $\Delta(y)$ be anything in $\mathcal{F}(y)$). Now look at the performance $\|x - \Delta\Phi(x)\|_X$ for some $x \in K$. Both x and $\Delta\Phi(x) =: x'$ are elements of K , so $x - x'$ is in $\mathcal{N}(\Phi)$ and in CK . Therefore $\frac{x - x'}{C} \in K \cap \mathcal{N}(\Phi)$. Thus,

$$\left\| \frac{x - x'}{C} \right\|_X \leq \sup_{z \in Y \cap K} \|z\|_X, \quad (3.9)$$

and so for any $x \in K$,

$$\|x - \Delta\Phi(x)\|_X \leq C \sup_{z \in Y \cap K} \|z\|_X. \quad (3.10)$$

Taking the infimum over all Y , we get that $E_n(K)_X \leq C d^n(K)_X$.

From the proof of this theorem, we see that there is a matching between the matrices Φ and the spaces Y (via the nullspace of Φ).

An important result is that $d^n(U(\ell_q^N))_{\ell_p^N}$ is known for all p, q except $p = 1, q = \infty$. A precise statement of these widths can be found in the book . A particularly important case is

$$C_1 \sqrt{\frac{\log(N/n)}{n}} \leq d^n(U(\ell_1^N))_{\ell_2^N} \leq C_2 \sqrt{\frac{\log(N/n)}{n}} \quad (3.11)$$

for $N > 2n$. This result was first proved by Kashin with a worse logarithm in the upper inequality and later brought to the present form by Gluskin. This result solves several important problems in functional analysis and approximation.

3.3 Instance Optimality⁵

Now we consider another way (actually two related ways) to measure optimality of an encoder/decoder pair.

⁵This content is available online at <<http://cnx.org/content/m15135/1.1/>>.

1. Instance optimality. Suppose we are in \mathbb{R}^N with an $n \times N$ measurement matrix Φ and a decoder Δ . Recall that

$$\sigma_k(x)_X := \inf_{z: \#z \leq k} \|x - z\|_X. \tag{3.12}$$

We say that the encoding/decoding strategy Φ, Δ is instance optimal of order k with constant C_0 if

$$\|x - \Delta\Phi(x)\|_X \leq C_0 \sigma_k(x)_X \tag{3.13}$$

for all $x \in \mathbb{R}^N$. (Note that we are no longer restricting x to a class K .) Better Φ 's have larger k for which this holds. The name "instance optimal" indicates that the encoding/decoding performance depends on each instance of x .

2. Mixed-norm instance optimality (MNIO). Let $q < p$. The encoder/decoder pair Φ, Δ is MNIO for p, q, k , and C_0 if

$$\|x - \Delta\Phi(x)\|_{\ell_p^N} \leq C_0 \frac{\sigma_k(x)_{\ell_q^N}}{k^{1/q-1/p}}. \tag{3.14}$$

Cases of interest include asking whether

$$\|x - \Delta\Phi(x)\|_{\ell_2^N} \leq C_0 \sigma_k(x)_{\ell_1^N}. \tag{3.15}$$

and whether

$$\|x - \Delta\Phi(x)\|_{\ell_2^N} \leq C_0 \frac{\sigma_k(x)_{\ell_1^N}}{\sqrt{k}}. \tag{3.16}$$

Let's focus on instance optimality. It would be interesting to know whether a given Φ satisfies this property. To answer this question, we state an equivalent condition to instance optimality.

Theorem 3.3:

Consider the statements

1. Φ, Δ is instance optimal of order k on X .
2. Φ has the following nullspace property (NSP): $\|\eta\|_X \leq C_1 \|\eta_{T^c}\|_X \forall \eta \in N(\Phi), \#T \leq k$.
3. $\|\eta\|_X \leq C_1 \sigma_k(\eta)_X \forall \eta \in N(\Phi)$.
4. $\|\eta_T\|_X \leq C'_1 \|\eta_T^c\|_X \forall \eta \in N(\Phi), \#T \leq k$.

Then (b) and (c) are equivalent with the same constant; (d) is equivalent to (b) and (c) but with a different constant. Also (a) with a value k implies (b) with the same k , and (b) with a value $2k$ implies (a) with a value k .

3.4 The Restricted Isometry Property⁶

We say that an $n \times N$ matrix Φ has the restricted isometry property (RIP) for k if for each $T \subseteq \{1, \dots, N\}$ such that $\#T \leq k$, Φ_T (the matrix formed by choosing the columns of Φ whose indices are in T) has the property

$(1 - \delta_k) \ x_T\ _{\ell_2} \leq \ \Phi_T(x)\ _{\ell_2} \leq (1 + \delta_k) \ x_T\ _{\ell_2}$	(RIP)
--	-------

Table 3.1

where $0 < \delta_k < 1$. This useful definition is by Candes and Tao. The idea is that the embedding of a k -dimensional space in M -dimensional space almost preserves norm – like an isometry. Another way of looking at it is to consider the matrix $\Phi_T^t \Phi_T$, of size $k \times k$. This matrix is symmetric, positive definite, and its eigen-values are between $1 - \delta_k$ and $1 + \delta_k$.

I prefer the following modified condition (dubbed the MIRP), which is more convenient for mathematical analysis:

⁶This content is available online at <<http://cnx.org/content/m15149/1.1/>>.

$(c_1)^{-1} \ x_T\ _{\ell_2} \leq \ \Phi_T(x)\ _{\ell_2} \leq c_1 \ x_T\ _{\ell_2}$	(MRIP)
---	--------

Table 3.2

We can now state the following theorem.

Theorem 3.4:

If Φ satisfies MRIP for $2k$ then $\exists \Delta$ s.t. (Φ, Δ) is instance optimal for ℓ_1^N for K .

This shows that whenever we have a matrix Φ satisfying the MRIP for $2k$ then it will perform well on encoding vectors (at least in the sense of ℓ_1^N accuracy). The question is how can we construct measurement matrices with this property? We can construct Φ using Gaussian entries and then normalizing the columns.

Theorem 3.5:

\exists constant $c > 0$ s.t. if $k \leq c \frac{n}{\log(N/n)}$ then with high probability Φ satisfies RIP and MRIP for k .

Given N and n , the range of k in the above results reflects how accurately we can recover data. There is another constant c' that serves as a converse bound for Theorem 3. This converse can be derived using Gluskin widths.

Rule 3.1:

The following generic problem is of great interest: Consider the class of matrices $\mathcal{M} = \{\Phi M \times N, \Phi \text{ has some prescribed property (eg. Toeplitz, circulant, etc.)}\}$. What is the largest k for which such a matrix can have the MRIP.

3.5 The Nullspace Property⁷

We begin with a property of the null space N which is at the heart of proving results on instance-optimality.

We say that N has the Null Space Property if for all $\eta \in N$ and all T with $\#T \leq k$ we have $\|\eta\|_X \leq c_1 \|\eta_{T^c}\|_X$

Intuitively, NSP implies that for any vector in the nullspace the energy will not be concentrated in a small number of entries.

The following are equivalent formulations for NSP X for k :

1. $\|\eta\|_X \leq c_1 \sigma_k(\eta)$
2. $\|\eta_T\|_X \leq c'_1 \|\eta_{T^c}\|_X$ where $\eta = \eta_T + \eta_{T^c}$.

Note also that the triangle inequality can be used as follows

$$\|\eta\|_X = \|\eta_T + \eta_{T^c}\|_X \leq \|\eta_T\|_X + \|\eta_{T^c}\|_X$$

which shows that (b) is equivalent to NSP.

Theorem 3.6:

1. If (Φ, Δ) is instance optimal on X for the value k , then Φ satisfies the NSP for $2k$ on X with an equivalent constant.
2. If Φ has the NSP for X and $2k$ then $\exists \Delta$ s.t. Φ has the instance optimal property for k .

⁷This content is available online at <http://cnx.org/content/m15148/1.1/>.

Proof:

We will prove a slightly weaker version of this to save time. We first prove that instance optimality for k implies NSP X for k (hence this is slightly weaker than advertised) . Let $\eta \in \mathcal{N}$ and set $z = \Delta(0)$ then

$$\begin{aligned} \|\eta - z\| &\leq c_0 \sigma_k(\eta) && \text{instance optimal property} \\ \|\eta + z\| &\leq c_0 \sigma_k(\eta) && -z \in \mathcal{N} \\ \|\eta\| &\leq \max\{\|\eta - z\|, \|\eta + z\|\} \leq c_0 \sigma_k(\eta) && \text{triangle inequality} \end{aligned} \quad (3.17)$$

We now prove 2. Suppose Φ has the NSP for $2k$. Given y , $\mathcal{F}(y) = \{x : \Phi(x) = y\}$. Let us define the decoder Δ by $\Delta(y) := \operatorname{argmin}\{\sigma_K(x)_X : x \in \mathcal{F}(y)\}$, then

$$\begin{aligned} \|x - \Delta(\Phi(x))\|_X &= \|x - x'\|_X \\ &\leq c_1 \sigma_{2K}(x - x') \\ &\leq c_1 (\sigma_K(x) - \sigma_K(x')) && \text{specific } 2K \text{ term approximation} \\ &\leq 2c_1 \sigma_K(x) \\ \text{QED.} \end{aligned}$$

Note that the instance optimal property automatically gives reproduction of K -sparse signals.

At this stage the challenge is to create Φ with this instance optimal property. For this we shall use the restricted isometry property as introduced earlier and which we now recall.

3.6 Optimality and the MRIP⁸

From our last lecture, we are interested in signals $x \in \mathbb{R}^N$, and we can make n measurements (or ask n questions) to obtain $y = \Phi x \in \mathbb{R}^n$. We proposed several optimality criteria to make these measurements, i.e., to choose the measurement matrix Φ .

1. For signals $x \in \Sigma_k$ (with k non-zero coefficients), we choose an encoder Φ and the corresponding decoder, Δ , such that $\Delta(\Phi(x)) = x$.
2. For classes of signals e.g. $K = U(\ell_p^N)$, our performance measure is closely related to the Gelfand widths.
3. Instance Optimal: our encoder and decoder, Φ and Δ , should satisfy

$$\|x - \Delta\Phi(x)\|_{\ell_p} \leq C_0 \sigma_k(x)_{\ell_p}. \quad (3.18)$$

We see that criterion (3) implies (1) since $\sigma_k(x)_{\ell_p} = 0$ for $x \in \Sigma_k$. We also showed in the previous lecture that we have instance optimality for order k if and only if we have the Null Space Property (NSP) of order $2k$. As a reminder, NSP of order $2k$ means that $\forall \eta \in \mathcal{N}, \mathcal{N} = \{x : \Phi(x) = 0\}$, we have $\|\eta\|_{\ell_p} \leq C_1 \sigma_{2k}(\eta)_{\ell_p}$. In other words, elements of $\eta \in \mathcal{N}$ are not sparse, and they are all of approximately equal size (i.e., they do not concentrate their entries in $2k$ positions).

We mentioned that, in order to attain instant optimality, the $n \times N$ measurement matrix Φ should have the modified restricted isometry property (MRIP) of order m , i.e., when we choose any m columns of Φ to obtain a $n \times m$ sub-matrix Φ_T where $m \leq n$,

$$C_2^{-1} \|x_T\|_{\ell_2} \leq \|\Phi_T x_T\|_{\ell_2} \leq C_2 \|x_T\|_{\ell_2}. \quad (3.19)$$

⁸This content is available online at <http://cnx.org/content/m15141/1.1/>.

Theorem 3.7:

If Φ has MRIP for $m = 3k$, then Φ has NSP for ℓ_1 of order $2k$, and so Φ is instance optimal in ℓ_1 of order k , i.e.,

$$\|x - \Delta\Phi(x)\|_{\ell_1} \leq C\sigma_k(x)_{\ell_1}. \quad (3.20)$$

A related result is that under the same assumption,

$$\|x - \Delta\Phi(x)\|_{\ell_2} \leq \frac{C\sigma_k(x)_{\ell_1}}{\sqrt{k}}. \quad (3.21)$$

Proof:

To prove ((3.20)), we need to only show NSP for ℓ_1 and $2k$, i.e.,

$$\|\eta\|_{\ell_1} \leq C\sigma_{2k}(\eta)_{\ell_1}, \eta \in \mathcal{N}. \quad (3.22)$$

Given η , let T_0 be the set of indices of the $2k$ largest entries, T_1 be the set of indices of the k next largest entries, T_2 be the set of indices of the k next largest entries, and so on.

$$\begin{aligned} \eta_0 &= \eta_{T_0} + \eta_{T_1}, \\ \eta &= \eta_{T_0} + \eta_{T_1} + \dots + \eta_{T_s}, \\ \Phi(\eta) &= \Phi(\eta_{T_0} + \eta_{T_1} + \dots + \eta_{T_s}) = 0, \\ \Phi(\eta_0) &= -\Phi(\eta_{T_2} + \dots + \eta_{T_s}), \text{ by linearity} \\ &= -\sum_{j=2}^s \Phi(\eta_{T_j}). \end{aligned} \quad (3.23)$$

Therefore, we can estimate

$$\begin{aligned} \|\eta_0\|_{\ell_2} &\leq C_2\|\Phi(\eta_0)\|_{\ell_2}, \text{ by restricted isometry} \\ &\leq C_2\sum_{j=2}^s\|\Phi(\eta_{T_j})\|_{\ell_2}, \text{ by triangle inequality} \\ &\leq C_2^2\sum_{j=2}^s\|\eta_{T_j}\|_{\ell_2}, \text{ by restricted isometry.} \end{aligned} \quad (3.24)$$

Since, for $j \geq 2$, η_{T_j} is the best k -term approximation to $\eta_{T_{j-1}} + \eta_{T_j}$,

$$\|\eta_{T_j}\|_{\ell_2} = \sigma_k(\eta_{T_{j-1}} + \eta_{T_j})_{\ell_2}. \quad (3.25)$$

Furthermore, we know for any $q < p$,

$$\sigma_k(\eta)_{\ell_p} \leq k^{\frac{1}{p} - \frac{1}{q}}\|x\|_{\ell_q}. \quad (3.26)$$

Combining ((3.25)) and ((3.26)), we obtain

$$\|\eta_{T_j}\|_{\ell_2} \leq \frac{\|\eta_{T_{j-1}} + \eta_{T_j}\|_{\ell_1}}{\sqrt{k}}. \quad (3.27)$$

Substituting ((3.27)) back into ((3.24)), we now have

$$\begin{aligned} \|\eta_0\|_{\ell_2} &\leq \frac{C_2^2}{\sqrt{k}}\sum_{j=2}^s\|\eta_{T_{j-1}} + \eta_{T_j}\|_{\ell_1} \\ &\leq \frac{C_2^2}{\sqrt{k}}\sum_{j=2}^s\|\eta_{T_{j-1}}\|_{\ell_1} + \|\eta_{T_j}\|_{\ell_1} \\ &\leq \frac{2C_2^2}{\sqrt{k}}\sum_{j=1}^s\|\eta_{T_j}\|_{\ell_1} \\ &\leq \frac{2C_2^2}{\sqrt{k}}\sigma_{2k}(\eta)_{\ell_1}. \end{aligned} \quad (3.28)$$

The last step is due to the fact that $\sum_{j=1}^s \|\eta_{T_j}\|_{\ell_1}$ is the $2k$ -term approximation error for η . Notice this is only true for ℓ_1 . This completes the proof of ((3.21)).

To prove ((3.20)), we let x_j be the j -th entry in η_{T_0} . By the Cauchy-Schwarz Inequality,

$$\begin{aligned} \|\eta_{T_0}\|_{\ell_1} &= \sum_{j=1}^{2k} |x_j| \\ &\leq \left(\sum_{j=1}^{2k} 1^2\right)^{\frac{1}{2}} \left(\sum_{j=1}^{2k} |x_j|^2\right)^{\frac{1}{2}}, \text{ by CSI} \\ &= \sqrt{2k} \|\eta_{T_0}\|_{\ell_2} \\ &\leq \sqrt{2k} \|\eta_0\|_{\ell_2} \\ &\leq 2\sqrt{2} C_2^2 \sigma_{2k}(\eta)_{\ell_1}, \end{aligned} \tag{3.29}$$

The $2k$ -term approximation error in ℓ_1 for η can be expressed as

$$\|\eta_{T_1} + \dots + \eta_{T_s}\|_{\ell_1} = \sigma_{2k}(\eta)_{\ell_1}. \tag{3.30}$$

Since

$$\eta = \eta_{T_0} + (\eta_{T_1} + \dots + \eta_{T_s}), \tag{3.31}$$

we can finally prove that

$$\begin{aligned} \|\eta\|_{\ell_1} &\leq \|\eta_{T_0}\|_{\ell_1} + \|\eta_{T_1} + \dots + \eta_{T_s}\|_{\ell_1}, \text{ by triangle inequality} \\ &= 2\sqrt{2} C_2^2 \sigma_{2k}(\eta)_{\ell_1} + \sigma_{2k}(\eta)_{\ell_1} \\ &= (2\sqrt{2} C_2^2 + 1) \sigma_{2k}(\eta)_{\ell_1}. \end{aligned} \tag{3.32}$$

Therefore, we have proved () with $C = 2\sqrt{2} C_2^2 + 1$. \square

3.7 Summary⁹

⁹This content is available online at <<http://cnx.org/content/m15145/1.1/>>.

3.7.1 Review

Last time we proved that for each $k \leq c_0 \frac{n}{\log N/n}$, there exists an $n \times N$ matrix Φ and a decoder Δ such that

- (a) - $\|x - \Delta\Phi(x)\|_{\ell_1} \leq c_0 \sigma_k(x)_{\ell_1}$
- (b) - $\|x - \Delta\Phi(x)\|_{\ell_2} \leq c_0 \frac{\sigma_k(x)_{\ell_1}}{\sqrt{k}}$

Recall that we can find such a Φ by setting the entries $[\Phi]_{j,k} = \phi_{j,k}(\omega)$ to be realizations of independent and identically distributed Gaussian random variables.

3.7.2 Deficiencies

3.7.2.1 Decoding is not implementable

Our decoding “algorithm” is:

$$\Delta(y) := \operatorname{argmin}_{x \in \mathcal{F}(y)} \sigma_k(x)_{\ell_1} \quad (3.33)$$

where $\mathcal{F}(y) := \{x : \Phi(x) = y\}$. In general, this algorithm is not implementable. This deficiency, however, is easily repaired. Specifically, define

$$\Delta_1(y) := \operatorname{argmin}_{x \in \mathcal{F}(y)} \|x\|_{\ell_1}. \quad (3.34)$$

Then (a) and (b) hold for Δ_1 in place of Δ . This decoding algorithm is equivalent to solving a linear programming problem, thus it is tractable and can be solved using techniques such as the interior point method or the simplex method. In general, these algorithms have computational complexity $O(N^3)$. For very large signals this can become prohibitive, and hence there has been a considerable amount of research in faster decoders (such as decoding using greedy algorithms).

3.7.2.2 We cannot generate such Φ

The construction of a Φ from realizations of Gaussian random variables is guaranteed to work with high probability. However, we would like to know, given a particular instance of Φ , do (a) and (b) still hold. Unfortunately, this is impossible to check (since, to show that Φ satisfies the MRIP for k , we need to consider all possible submatrices of Φ). Furthermore, we would like to build Φ that can be implemented in circuits. We also might want fast decoders Δ for these Φ . Thus we also may need to be more restrictive in building Φ . Two possible approaches that move in this direction are as follows:

1. Find Φ that we can build such that we can prove instance optimality in ℓ_1 for a smaller range of k , i.e.,

$$\|x - \Delta\Phi(x)\|_{\ell_1} \leq c_0 \sigma_k(x)_{\ell_1} \quad (3.35)$$

for $k < K$. If we are willing to sacrifice and let K be smaller than before, for example, $K \approx \sqrt{n}$, then we might be able to prove that $\Phi_T^t \Phi_T$ is diagonally dominant for all T such that $\#T = 2k$, which would ensure that Φ satisfies the MRIP.

2. Consider $\Phi(\omega)$ where ω is a random seed that generates a Φ . It is possible to show that give x , with high probability, $\Phi(\omega)(x) = y$ encodes x in an ℓ_2 -instance optimal fashion:

$$\|x - \bar{x}\|_{\ell_2} \leq 2\sigma_k(x)_{\ell_2} \quad (3.36)$$

for $k \leq c_0 \frac{n}{(\log N/n)^{5/2}}$. Thus, by generating many such matrices we can recover any x with high probability.

3.7.2.3 Encoding signals

Another practical problem is that of encoding the measurements y . In a real system these measurements must be quantized. This problem was addressed by Candes, Romberg, and Tao in their paper *Stable Signal Recovery from Incomplete and Inaccurate Measurements*. They prove that if y is quantized to \bar{y} , and if $x \in U(\ell_p)$ for $p \leq 1$, then we get optimal performance in terms the number of bits required for a given accuracy. Notice that their result applies only to the case where $p \leq 1$. One might expect that this argument could be extended to p between 1 and 2, but a warning is in order at this stage:

Fix $1 < p \leq 2$. Then there exist Φ and Δ satisfying

$$\|x - \Delta\Phi(x)\|_{\ell_p} \leq C_0\sigma_k(x)_{\ell_p} \quad (3.37)$$

if

$$k \leq c_0 N^{\frac{2-2/p}{1-2/p}} \left(\frac{n}{\log N/n} \right)^{\frac{p}{2-p}}. \quad (3.38)$$

Furthermore, this range of k is the best possible (save for the \log term).

Examples:

- $p = 1$, we get our original results
- $p = 2$, we do not get instance optimal for $k = 1$ unless $n \approx N$
- $p = \frac{3}{2}$, we only get instance optimal if $k \leq c_0 N^{-2} \left(\frac{n}{\log N/n} \right)^3$

Attributions

Collection: *Compressive Sensing*

Edited by: Mark A. Davenport, Richard Baraniuk, Ronald DeVore

URL: <http://cnx.org/content/col10458/1.1/>

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Introduction"

By: Ronald DeVore

URL: <http://cnx.org/content/m15136/1.1/>

Page: 1

Copyright: Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "The Shannon-Whitaker Sampling Theorem"

By: Michael A Lexa, Ronald DeVore

URL: <http://cnx.org/content/m15146/1.2/>

Pages: 3-4

Copyright: Michael A Lexa, Ronald DeVore

License: <http://creativecommons.org/licenses/by/3.0/>

Module: "Optimal Encoding"

By: Ronald DeVore, Shriram Sarvotham

URL: <http://cnx.org/content/m15139/1.1/>

Page: 5

Copyright: Ronald DeVore, Shriram Sarvotham

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Kolmogorov Entropy"

By: Ronald DeVore, Shriram Sarvotham

URL: <http://cnx.org/content/m15137/1.1/>

Page: 6

Copyright: Ronald DeVore, Shriram Sarvotham

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Optimal Encoding of Bandlimited Signals"

By: Ronald DeVore, Shriram Sarvotham

URL: <http://cnx.org/content/m15140/1.2/>

Pages: 7-8

Copyright: Ronald DeVore, Shriram Sarvotham

License: <http://creativecommons.org/licenses/by/3.0/>

Module: "Stable Signal Representations"

By: Michael A Lexa, Ronald DeVore, Shriram Sarvotham

URL: <http://cnx.org/content/m15144/1.1/>

Pages: 9-11

Copyright: Michael A Lexa, Ronald DeVore, Shriram Sarvotham

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Preliminaries"

By: Marco F. Duarte, Ronald DeVore

URL: <http://cnx.org/content/m15142/1.1/>

Page: 13

Copyright: Marco F. Duarte, Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "New Signal Models"

By: Mark A. Davenport, Ronald DeVore

URL: <http://cnx.org/content/m15138/1.1/>

Pages: 13-15

Copyright: Mark A. Davenport, Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Sparse Approximation and ℓ_p Spaces"

By: Marco F. Duarte, Ronald DeVore

URL: <http://cnx.org/content/m15143/1.1/>

Pages: 15-17

Copyright: Marco F. Duarte, Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Thresholding and Greedy Bases"

By: Jason Laska, Ronald DeVore

URL: <http://cnx.org/content/m15147/1.1/>

Pages: 17-19

Copyright: Jason Laska, Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Greedy Algorithms"

By: Jason Laska, Ronald DeVore

URL: <http://cnx.org/content/m15134/1.1/>

Pages: 19-21

Copyright: Jason Laska, Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Compressive Sensing"

By: Mark A. Davenport, Ronald DeVore, Chris Rozell

URL: <http://cnx.org/content/m15132/1.1/>

Pages: 23-24

Copyright: Mark A. Davenport, Ronald DeVore, Chris Rozell

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Gelfand n-widths"

By: Mark A. Davenport, Ronald DeVore, Chris Rozell, Michael Wakin

URL: <http://cnx.org/content/m15133/1.1/>

Pages: 25-26

Copyright: Mark A. Davenport, Ronald DeVore, Chris Rozell, Michael Wakin

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Instance Optimality"

By: Ronald DeVore

URL: <http://cnx.org/content/m15135/1.1/>

Pages: 26-27

Copyright: Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

ATTRIBUTIONS

37

Module: "The Restricted Isometry Property"

By: Ronald DeVore

URL: <http://cnx.org/content/m15149/1.1/>

Pages: 27-28

Copyright: Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "The Nullspace Property"

By: Ronald DeVore

URL: <http://cnx.org/content/m15148/1.1/>

Pages: 28-29

Copyright: Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Optimality and the MRIP"

By: Wai Lam Chan, Ronald DeVore

URL: <http://cnx.org/content/m15141/1.1/>

Pages: 29-31

Copyright: Wai Lam Chan, Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

Module: "Summary"

By: Mark A. Davenport, Ronald DeVore

URL: <http://cnx.org/content/m15145/1.1/>

Pages: 32-34

Copyright: Mark A. Davenport, Ronald DeVore

License: <http://creativecommons.org/licenses/by/2.0/>

About Connexions

Since 1999, Connexions has been pioneering a global system where anyone can create course materials and make them fully accessible and easily reusable free of charge. We are a Web-based authoring, teaching and learning environment open to anyone interested in education, including students, teachers, professors and lifelong learners. We connect ideas and facilitate educational communities.

Connexions's modular, interactive courses are in use worldwide by universities, community colleges, K-12 schools, distance learners, and lifelong learners. Connexions materials are in many languages, including English, Spanish, Chinese, Japanese, Italian, Vietnamese, French, Portuguese, and Thai. Connexions is part of an exciting new information distribution system that allows for **Print on Demand Books**. Connexions has partnered with innovative on-demand publisher QOOP to accelerate the delivery of printed course materials and textbooks into classrooms worldwide at lower prices than traditional academic publishers.