SAMPLING AND RECONSTRUCTION IN THE 21ST CENTURY

Justin Romberg

Georgia Tech, School of Electrical and Computer Engineering

ABSTRACT

We review advances in sampling theory since the turn of the century, with a special emphasis on how basis decompositions have allowed us to pose and analyze sampling problems more through the lens of linear algebra, than through the traditional lens of filtering and Fourier decompositions. This new framework also allows us to incorporate nonlinear (sparse) signal models into the signal reconstruction process using optimization.

Index Terms— Sampling, splines, compressed sensing

1. INTRODUCTION

In the past 15 years, our understanding of what it means to sample or digitize a signal has broadened. Much of this progress is owed to new interfaces that have opened up between signal processing and applied mathematics. This paper discusses the influence of two areas in particular, harmonic analysis, which gives us a more general notion of what it means to discretize a continuous time signal, and optimization, which makes it possible to incorporate nonlinear modeling techniques into the reconstruction process.

Our discussion will be informal, and while there are some key references given, the text below is meant more to reflect the author's viewpoint of how these topics fit together, rather than serve as a comprehensive review. In Section 2, we start by discussing how the success of the wavelet transform made two important concepts mainstream: there are ways to decompose signals that are not based on sinusoids, and that simple nonlinear models based on sparsity can be very powerful. In Section 3, we discuss how the abstract concept of decomposing a signal in a basis has affected the way we think about traditional sampling architectures. In Sections 4 and 5, we discuss the closely related problems of finding the sparsest decomposition in an overcomplete representation, and recovering a signal from a small number of linear measurements. These problems have created a central place for optimization in signal processing, just as wavelets introduced basic functional analysis to a broad audience of engineers.

Email: jrom@ece.gatech.edu. This work was supported by NSF grant CCF-1422540.

2. BASES AND SPARSITY

In the 1990s, applications in signal and image processing based on the wavelet transform flourished. Their success hinged on two qualitative properties: they were local, meaning that sharp events in the signal only impacted a small number of coefficients, and stable, in that small changes in the wavelet coefficients led to comparably small changes in the signal. Wavelets are particularly well-suited for analyzing signals that are non-stationary in that they have sharp transients; these transients are reflected in the transform coefficients in a way that is far more obvious than it is in the Fourier domain. Their design and implementation were also well-suited to the strengths of signal processing researchers. The wavelet transform can be interpreted as a time-frequency decomposition, and its structure allows it to be implemented as a series of digital filters and re-sampling devices.

Along with state-of-the-art noise removal [1] and compression [2, 3] algorithms, there were deep theoretical results which demonstrated that very simple algorithms applied to the wavelet coefficients of signals that were piecewise smooth yield results that are asymptotically optimal (overviews of this theory can be found in the review paper [4] and the book [5]).

Along with these advances in concrete applications and theory, the wavelet transform also had a significant philosophical impact in the signal and image processing community. On a high level, the research on wavelet demonstrated that there was true value in decomposing a signal using something other than sinusoids. The general concept of an atomic decomposition, taking a continuous-time signal $x(t)$ and writing it as

$$x(t) = \sum_n \alpha_n \psi_n(t), \tag{1}$$

where the $\{\psi_n(t)\}$ were a fixed set of basis functions that might vary depending on what kinds of signals you were interested in, had become mainstream in signal processing. Mapping the continuous-time $x(t)$ to the discrete sequence of numbers $\{\alpha_n\}$ is simply a way to discretize $x(t)$, the utility of which is clear if you are going to process the signal digitally.

If we are using the coefficient sequence $\{\alpha_n\}$ to represent $x(t)$, it is important that the mappings from the signal to the expansion coefficients and back be stable. This ensures that each $\boldsymbol{x}$ maps to a different $\{\alpha_n\}$, and that perturbing the coefficients slightly will not lead to massive changes during the

reconstruction process. These conditions are ensured with the following condition on the linear operator $\mathcal{B}$ that maps signals $x(t)$ to their coefficients $\boldsymbol{\alpha}$:

$$A\|\boldsymbol{x}\|_2^2 \leq \|\mathcal{B}(\boldsymbol{x})\|_2^2 \leq B\|\boldsymbol{x}\|_2^2. \quad (2)$$

The norm in the middle of the expression above is the sum-of-squares of the coefficient sequence $\{\alpha_n\}$ in (1), while the $\|\boldsymbol{x}\|_2^2$ on the left and right denotes the energy in the continuous-time signal $x(t)$.

The wavelet transforms also came with a new modeling paradigm. A very clear mathematical explanation emerged for why using wavelets for the ψ_n in (1) would work better than sinusoids for certain kinds of signals: wavelets were more effective because they used fewer terms to represent the signal. That is, for signals of interest, most of the α_n in (1) are so small as to be insignificant. This means that an accurate, low-order approximation to $x(t)$ can be formed by cherry-picking the largest terms[1].

Despite its simplicity, this *sparsity* model is highly nonlinear. It says that we can approximate the signal as living in a low-dimensional subspace, but this choice of subspace depends on the signal itself. More to the point, the significant coefficients can be different for every signal — the wavelet transform is local by design, and as the transients in the signal move around so will the locations of the significant wavelet coefficients.

3. SAMPLING

Basis and frame expansions have also helped us generalize the most fundamental result in all of signal processing: the Shannon-Nyquist sampling theorem.

All of digital signal processing relies on our ability to take a continuous-time signal and capture it with a discrete list of numbers. The classical Shannon-Nyquist theorem tell us that this can be done in the most straightforward way possible: if the signal is bandlimited, then we will be able to recover it from equally spaced samples provided that those samples are taken at a rate that is at least twice the highest frequency in the signal. Understanding this result and its implications is the main event in undergraduate signal processing courses. Along with being a fundamental result, it is a crowd-pleaser: there is something magical about being able to fill in the gaps between discrete samples with confidence that what you are doing is exactly right.

3.1. Sampling splines

Shannon-Nyquist can be, and usually is, understood fully without connecting it to a basis expansion of the form (1). But the extensions to the theory developed around the turn of the century, which were also based on the physically-realizable architecture of filter-then-sample, rely critically on this interpretation. One extension of the classical theory, reviewed carefully in [6], shows how polynomial splines, another important class of signals, can be captured with equally spaced samples. It was well-known in the numerical analysis and approximation theory communities that pth order polynomial splines, which are functions that are pth order polynomials between equally spaced knots, can be written as a superposition of shifted B-spline functions. In (1), we take

$$\psi_n(t) = \psi(t - nT), \quad \text{where} \quad \hat{\psi}(\omega) = \left(\frac{\sin(\omega/2)}{\omega/2}\right)^{p+1}.$$

Since the basis functions in the synthesis formula are shifted versions of the template $\psi(t)$, the synthesis of $x(t)$ from its basis coefficients $\{\alpha_n\}$ can be implemented using the same type of architecture as reconstructing a bandlimited signal from its point-samples simply by replacing the ideal lowpass filter in the latter with a filter that has a frequency response equal to $\hat{\psi}(\omega)$ above. Unlike bandlimited reconstruction, the atoms $\psi_n(t)$ for spline reconstruction are strictly local in time, meaning that the influence of a particular coefficient α_{n_0} only effects the synthesized signal over a time interval of length $p + 1$.

The connection to classical sampling goes even deeper. The basis coefficients $\{\alpha_n\}$ can themselves be computed using a filter-then-sample architecture. The B-spline basis is stable, in that (2) holds for A and B reasonably close to one another, and so there exists a *dual basis* $\{\tilde{\psi}_n(t)\}$ such that the expansion coefficients can be computed as

$$\alpha_n = \langle \boldsymbol{x}, \tilde{\boldsymbol{\psi}}_n \rangle. \quad (3)$$

It happens that since the ψ_n are all shifts of the same function, the same will be true for the dual functions, $\tilde{\psi}_n = \tilde{\psi}(t - nT)$ for some $\tilde{\psi}(t)$ that we can compute. This means that the series of inner products we need to compute the expansion coefficients in (3) are (equally spaced) samples of the convolution of $\tilde{\psi}(t)$ and $x(t)$.

We now see all of the parallels between the acquisition strategy for bandlimited signals and that for polynomial splines. We sample a bandlimited signal by passing it through an ideal lowpass filter (the anti-aliasing filter), then taking uniform samples. We compute the basis expansion coefficients (which might be thought of as generalized samples) of a polynomial spline by passing it through a filter with frequency response $\hat{\tilde{\psi}}(\omega)$, then taking uniform samples. Given the samples/coefficients, the reconstruction process is also very similar: the discrete sequence is converted into a train of weighted Dirac delta functions, then passed through an analog filter; this filter is an ideal lowpass filter in the bandlimited case, and has a frequency response of $\hat{\psi}(\omega)$ in the polynomial spline case.

The lesson here is that acquisition strategies of the general form of "filter then sample uniformly" can be applied to

[1] That the largest coefficients are actually the most important is exactly true if the $\psi_n(t)$ are orthonormal; if the basis $\{\psi_n\}$ is well-conditioned, then choosing the largest is almost optimal.

many different kinds of signal models. What is important is not that the signal is bandlimited, but rather that it can be written (or at least approximated) in a stable way as the superposition of equally spaced basis functions. Even though the acquisition and reconstruction operators consist of basic signal processing building blocks (filters, sampling devices), the development and our understanding of these generalizations would not be possible without looking at them through the lens of a basis decomposition.

3.2. Finite rates of innovation

In both the bandlimited and polynomial spline sampling results described above, our ability to reconstruct the signal comes from the underlying assumption that it lives in a known subspace. This is a linear model; the signal reconstruction techniques can be interpreted as solving a system of linear equations (perhaps with an arbitrarily large number of variables), and the entire acquisition and reconstruction pipeline can be interpreted as a projection of the input signal onto this subspace.

Around 2002, researchers started to ask whether the same models (i.e. sparsity) that sat at the core of compression and restoration algorithms might also play a role in sampling theory. In [7], the idea of sampling a signal with a "finite rate of innovation" was introduced; the model here was not that a signal could be built up out of a fixed library of basis functions, but rather that it contained only a limited number of highly localized events over a certain period of time.

The most straightforward example of such a model is to assume that the signal is a periodic stream of Dirac delta functions; that is, we write

$$x(t) = \sum_{k=1}^{K} a_k \delta(t - t_k), \quad \text{for} \quad t, t_k \in [0, 1], \tag{4}$$

and then copy the fundamental interval to make $x(t)$ periodic over the entire real line. It is important to stress that both the Dirac weights a_k and locations t_k are unknown. To acquire the signal, it is passed though a lowpass filter, and then sampled uniformly at M locations on $[0, 1]$ (the passband of the filter is matched to M). The initial filtering has the effect of taking the Fourier series of the original signal and limiting it to M terms — these M Fourier terms are readily calculated given the equally spaced samples.

Estimating the $\{a_k, t_k\}$ is now reduced to a very classical problem. The $\delta(t - t_k)$ express themselves as sinusoids in the Fourier domain, so the M Fourier coefficients are M equally spaced samples of superposition of K sinusoids with unknown frequency and unknown complex amplitude. There are many established methods from spectral analysis for attacking this problem. (We note also the connections to the recent work in super-resolution in [8].) This technique is also extensible to models other than the superposition of delta functions in (4). For example, signals that are piecewise polynomial with knots at unknown locations can also be acquired using similar ideas.

Here, the acquisition framework is very similar to the classical techniques: filter then sample uniformly. The estimation procedure, however, is completely different, and is highly nonlinear.

4. SPARSE EXPANSIONS

In parallel to the developments in sampling, sparsity was starting to play a role in computing the expansion coefficients for *overcomplete* linear decompositions. While having a stable representation of the form (2) means that $\mathcal{B}(\boldsymbol{x})$ is different for every signal $\boldsymbol{x}$ of interest, it does not necessarily mean that there is only one way to write each $\boldsymbol{x}$ in the form (1). It is often natural to use sets of signals $\{\psi_n\}$ that are linearly dependent; a simple example would be if a bandlimited signal (or a polynomial spline, for that matter) were sampled at a rate higher than necessary for perfect reconstruction. A linear representation which is stable in the sense of (2) but has elements that are linearly dependent is called a *frame*[9]. Processing frame coefficients is often times more robust than processing the coefficients of a critically sampled decomposition, as the redundancy causes errors to average out during reconstruction.

With an overcomplete frame, there are many ways a signal can be represented. The natural way to do this, using the set of coefficients that has minimum energy, leads to a frustrating result: even if a signal is indeed a superposition of a small number of frame elements, computing the frame expansion typically leads to many (if not all) non-zero coefficients. This is best illustrated with a simple, finite dimensional example. Suppose that a $\boldsymbol{x} \in \mathbb{C}^N$ can be written as a superposition of a small number of (discrete) spikes and a small number of (discrete) sinusoids. This means that there is a sparse vector $\boldsymbol{\alpha}_0$ such that

$$\boldsymbol{x} = \begin{bmatrix} \mathbf{I} & \boldsymbol{F} \end{bmatrix} \boldsymbol{\alpha}_0, \tag{5}$$

where $\mathbf{I}$ is the identity (spike basis), and $\boldsymbol{F}$ is the standard discrete Fourier matrix (sinusoidal basis). The minimal energy coefficients that reproduce a given $\boldsymbol{x}$ are given by $\boldsymbol{x}/2$ and $\boldsymbol{F}^{\mathrm{H}}\boldsymbol{x}/2$ stacked upon one another — neither of these vectors is sparse, as the spike component of $\boldsymbol{F}$ will be spread out in frequency, while the sinusoidal component will be spread out in time.

If there is a sparse decomposition in a frame, how can you find it? There are multiple approaches. The first, and perhaps the easiest to understand, is that you formulate a greedy selection procedure to find frame elements that explain your signal [10, 11]. If you concatenate your frame elements as columns in a matrix $\boldsymbol{\Psi}$, these "matching pursuit" algorithms find the element most correlated with $\boldsymbol{x}$, subtract it out, and then repeat with the residual.

Another approach, first analyzed in [12] for the spike+sinusoid case in (5), it to set up a global optimization program to find the sparsest explanation for $\boldsymbol{x}$. As stated, this is a hard problem to solve, but surprisingly using the ℓ_1 norm as a convex surrogate for sparsity turns out to be remarkably effective, both in theory and in practice. Given a signal $\boldsymbol{x}$, we decompose it in an overcomplete dictionary $\boldsymbol{\Psi}$ by solving

$$\underset{\boldsymbol{\alpha}}{\text{minimize}} \|\boldsymbol{\alpha}\|_1 \quad \text{subject to} \quad \boldsymbol{\Psi}\boldsymbol{\alpha} = \boldsymbol{x}. \tag{6}$$

Solving the program above is very tractable computationally, and if there is indeed a sparse enough solution, it is guaranteed to find it [13, 14].

5. COMPRESSIVE SAMPLING

The work on finding the sparsest representation in a frame was an example of a broader phenomenon that would soon become fully appreciated: it is possible to meaningfully solve underdetermined systems of equations if the solution is sparse. The effort to harness this idea to say something about data acquisition followed quickly; this body of work eventually became known as "compressed sensing" or "compressive sampling".

Compressive sampling (CS) works with the signal acquisition and reconstruction problem entirely in the framework of linear algebra. The computations in CS are all discrete, but the connection to sampling and reconstructing an analog signal is made in a familiar way: we start by assuming that the (continuous-time) signal $x(t)$ that we are interested in acquiring can be written (or closely approximated) as a superposition of a finite number of continuous-time basis functions $\psi_n(t)$; meaning the sum in (1) runs over $n = 1, \ldots, N$. We observe M linear functionals of $\boldsymbol{x}$, possibly corrupted by noise

$$y_m = \mathcal{L}_m(\boldsymbol{x}) + \text{noise}.$$

We can rewrite the measurements in matrix form by plugging in the linear expansion for $\boldsymbol{x}$:

$$\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_M \end{bmatrix} = \begin{bmatrix} \mathcal{L}_1(\boldsymbol{\psi}_1) & \mathcal{L}_1(\boldsymbol{\psi}_2) & \cdots & \mathcal{L}_1(\boldsymbol{\psi}_N) \\ \mathcal{L}_2(\boldsymbol{\psi}_1) & \mathcal{L}_2(\boldsymbol{\psi}_2) & \cdots & \mathcal{L}_2(\boldsymbol{\psi}_N) \\ \vdots & & \ddots & \\ \mathcal{L}_M(\boldsymbol{\psi}_1) & \mathcal{L}_M(\boldsymbol{\psi}_2) & \cdots & \mathcal{L}_M(\boldsymbol{\psi}_N) \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_N \end{bmatrix}$$

The formulation is general in that the $\mathcal{L}_m(\cdot)$ can be anything (e.g. point-sampling operators, samples of convolutions with a known kernel, or inner products against fixed test functions) as long as they are linear.

The question, then, is when $\boldsymbol{\alpha}$ can be stably recovered from $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{\alpha}$. Using the linear model, where all we know about $\boldsymbol{x}$ is that it is in the span of the $\{\boldsymbol{\psi}_n\}$, the answer is classical: the recovery is stable when $\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}$ is well-conditioned. This is equivalent to the frame bound in (2) with the matrix $\boldsymbol{\Phi}$ replacing the linear operator $\mathcal{B}$.

When the signal is sparse, meaning that only a small number of the entries in $\boldsymbol{\alpha}$ are significant, then recovery is possible under much weaker conditions. In [15], it was shown that it is sufficient for $\boldsymbol{\Phi}$ to behave as a *stable embedding* for sparse vectors. More precisely, if for some sufficiently small $\delta < 1$,

$$(1-\delta)\|\boldsymbol{v}_1 - \boldsymbol{v}_2\|_2^2 \leq \|\boldsymbol{\Phi}(\boldsymbol{v}_1 - \boldsymbol{v}_2)\|_2^2 \leq (1+\delta)\|\boldsymbol{v}_1 - \boldsymbol{v}_2\|_2^2,$$

for all $\boldsymbol{v}_1, \boldsymbol{v}_2$ that are S-sparse, then we can recover an S-sparse vector $\boldsymbol{\alpha}$ from $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{\alpha}$. There are a number of methods that produce an exact recovery when the measurements are not noisy, and a stable approximation when the measurements contain noise. A very thorough review of the theoretical results in this field can be found in [16].

It is now natural to ask what kind of $\boldsymbol{\Phi}$ provide this stable embedding for sparse vectors using an "efficient" number of rows. The answer to this question provides an interesting philosophical twist. In short, the the best $\boldsymbol{\Phi}$ are diverse, with each measurement mixing together a significant percentage of the expansion coefficients in different ways. This means that rather than being finely tuned to the basis structure of the signal, the measurement operators $\mathcal{L}_m(\cdot)$ should "touch" almost all of the basis functions for every m. In fact, if the entries of $\boldsymbol{\Phi}$ are generated at random, then S-sparse vectors can be recovered from

$$M \geq \text{Const} \cdot S \log(N/S)$$

measurements. That is, the number of rows we need in $\boldsymbol{\Phi}$ to solve $\boldsymbol{\Phi}\boldsymbol{\alpha} = \boldsymbol{y}$ in a meaningful way scales not with the number of columns, but rather with the number of non-zero terms we expect in the solution. This idea that diversity can be achieved by injecting randomness into the acquisition process, and that this diversity can pay real dividends when the time comes to estimate the signal, has had influence even outside the field of CS (see, for example, [17] and [18]).

One way we might interpret solving an optimization program like (6) is that we are using a model to help regularize an ill-conditioned inverse problem. The idea of introducing a statistical prior, forward measurement model, and then solving Bayes equation to perform the recovery is commonplace, especially in imaging, both inside and outside the signal processing community. But the effectiveness of recovering a sparse signal using a convex program is really best understood in a different way. Programs like (6) are not effective because sparse signals are "typical" vectors with ℓ_1 norm of a certain size; in fact, nearly all such vectors are not sparse at all. What is important is that sparse signals lie on the *singular facets* of the exterior of the set of all signals with comparable ℓ_1 norms. Connecting this type of geometrical picture to concrete mathematical guarantees is done beautifully in [19, 20].

6. REFERENCES

[1] R. Coifman and D. Donoho, "Translation-invariant de-noising," in *Wavelets and Statistics*, Anto-

niadis/Oppenheim, Ed., vol. 103 of *Lecture Notes in Statistics*. Springer-Verlag, 1995.

[2] A. Skodras, C. Christopoulos, and T. Ebrahimi, "The JPEG2000 still image compression standard," *IEEE Signal Proc. Mag.*, vol. 18, pp. 36–58, Sept 2001.

[3] Z. Xiong, K. Ramchandran, and M. T. Orchard, "Space-frequency quantization for wavelet image coding," *IEEE Trans. Image Proc.*, vol. 6, no. 5, pp. 677–693, 1997.

[4] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies, "Data compression and harmonic analysis," *IEEE Trans. Inform. Theory*, vol. 44, pp. 2435–2476, October 1998.

[5] S. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, Academic Press, 3rd edition, 2009.

[6] M. Unser, "Sampling — 50 years after Shannon," *Proceedings of the IEEE*, vol. 88, no. 4, pp. 569–587, April 2000.

[7] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Proc.*, vol. 50, pp. 1417–1428, June 2002.

[8] E. J. Candès and C. Fernandez-Granda, "Towards a mathematical theory of super-resolution," *Comm. Pure Appl. Math.*, vol. 67, no. 6, pp. 906–956, June 2014.

[9] J. Kovacevic and A. Chebira, "Life beyond bases: The advent of frames," *IEEE Signal Proc. Mag.*, July 2007.

[10] G. Davis, S. Mallat, and M. Avellaneda, "Adaptive greedy approximation," *J. Constructive Approx.*, vol. 12, pp. 57–98, 1997.

[11] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inform. Theory*, vol. 50, no. 10, pp. 2231–2242, October 2004.

[12] D. L. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inform. Theory*, vol. 47, no. 7, pp. 2845–2862, 2001.

[13] R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Trans. Inform. Theory*, vol. 49, pp. 3320–3325, 2003.

[14] M. Elad and A. M. Bruckstein, "A generalized uncertainty principle and sparse representation in paris of r^n bases," *IEEE Trans. Inform. Theory*, vol. 48, pp. 2558–2567, 2002.

[15] E. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5245, December 2006.

[16] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, Birkhauser, 2013.

[17] R. Raskar, A. Agrawal, and J. Tumblin, "Coded exposure photography: motion deblurring using fluttered shutter," in *Proc. ACM SIGGRAPH*, 2006.

[18] L. Carin, "On the relationship between compressive sensing and random sensor arrays," *IEEE Antennas and Propagation*, vol. 52, no. 5, pp. 72–81, 2009.

[19] D. L. Donoho and J. Tanner, "Neighborliness of randomly-projected simplices in high dimensions," *Proc. Natl. Acad. Sci. USA*, vol. 102, no. 27, pp. 9452–9457, 2005.

[20] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The convex geometry of linear inverse problems," *Found. of Comput. Math.*, vol. 12, no. 6, pp. 805–849, 2012.