



IRWIN AND JOAN JACOBS
CENTER FOR COMMUNICATION AND INFORMATION TECHNOLOGIES

Beyond Bandlimited Sampling: Nonlinearities, Smoothness and Sparsity

Y. C. Eldar and T. Michaeli

CCIT Report # 698
June 2008

■ ■ ■ ■ Electronics
■ ■ ■ ■ Computers
■ ■ ■ ■ Communications

DEPARTMENT OF ELECTRICAL ENGINEERING
TECHNION - ISRAEL INSTITUTE OF TECHNOLOGY, HAIFA 32000, ISRAEL



Beyond Bandlimited Sampling: Nonlinearities, Smoothness and Sparsity

Y. C. Eldar and T. Michaeli

Digital applications have developed rapidly over the last few decades. Since many sources of information are of analog or continuous-time nature, discrete-time signal processing (DSP) inherently relies on sampling a continuous-time signal to obtain a discrete-time representation. Consequently, sampling theories lie at the heart of signal processing applications and communication systems. A few examples are sampling rate conversion between audio formats and for software radio [1], biomedical imaging [2], lens distortion correction, super-resolution of image sequences and more.

To accommodate high operating rates while retaining low computational cost, efficient analog-to-digital (ADC) and digital-to-analog (DAC) converters must be developed. Many of the limitations encountered in current converters is due to a traditional assumption that the sampling stage must acquire the data at the Shannon-Nyquist rate, corresponding to twice the signal bandwidth [3], [4], [5]. To avoid aliasing, a sharp low-pass filter must be implemented prior to sampling. The reconstructed signal is also a bandlimited function, generated by integer shifts of the sinc interpolation kernel.

A major drawback of this paradigm is that many natural signals are better represented in alternative bases other than the Fourier basis [6], [7], [8], or possess further structure in the Fourier domain. In addition, ideal point-wise sampling, as assumed by the Shannon theorem, cannot be implemented. More practical ADCs introduce a distortion which should be accounted for in the reconstruction process. Finally, implementing the infinite sinc interpolating kernel is difficult, since it has slow decay. In practice, much simpler kernels are used such as linear interpolation. Therefore there is a need to develop a general sampling theory that will accommodate an extended class of signals beyond bandlimited functions, and will account for the nonideal nature of the sampling and reconstruction processes.

Sampling theory has benefited from a surge of research in recent years, due in part to the intense research in wavelet theory and the connections made between the two fields. In this survey we present several extensions of the Shannon theorem, that have been developed primarily in the past two decades, which treat a wide class of input signals as well as nonideal sampling and nonlinear distortions. This framework is based on viewing sampling in a broader sense of projection onto appropriate subspaces, and then choosing the subspaces to yield interesting

Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa 32000, Israel. E-mail: {yonina@ee,tomermic@tx}.technion.ac.il. Fax: +972-4-8295757. Tel: +972-4-8293256.

new possibilities such as pointwise sampling of non bandlimited signals, and perfect compensation of nonlinear effects.

Our focus here is on shift-invariant settings in which both sampling and reconstruction are obtained by filtering operations, and the sampling grid is uniform. However, all the results herein can be extended to arbitrary Hilbert space settings [9], [10], [11], [12] including finite-dimensional spaces, spaces that are not shift invariant and nonuniform sampling. Our exposition is based on a Hilbert-space interpretation of sampling techniques, and relies on the concepts of bases and projections. This perspective has been motivated in the context of sampling in the excellent review by Unser [8]. Here we consider a similar setting and compliment the paper of Unser by surveying further progress made in this area in recent years.

We begin by presenting a broad class of sampling theorems for signals confined to an arbitrary subspace in the presence of non-ideal sampling, and nonlinear distortions. Surprisingly, many types of nonlinearities that are encountered in practice do not possess any technical difficulty and can be completely compensated for. Next, we develop minimax recovery techniques that best approximate an arbitrary smooth input signal. These methods can also be used to reconstruct a signal using a given interpolation kernel that is easy to implement, with only a minor loss in signal quality. To further enhance the quality of the interpolated signal, we discuss fine grid recovery techniques in which the system rate is increased during reconstruction. As we show, the algorithms we develop can all be extended quite naturally to the recovery of random signals. These additional aspects extend the existing sampling framework and incorporate more realistic sampling and interpolation models.

Before proceeding with the detailed development, we note that another topic in the context of sampling that has gained growing attention recently is that of reconstructing signals that are known to be sparse in some domain. This class of problems is the underlying of the emerging field of compressed sensing [13], [14]. However, this framework has focused primarily on sampling of discrete signals and reconstruction techniques from a finite number of samples, while our interest here is on sampling and reconstructing analog continuous-time signals from uniform samples. Some exceptions are the work in [15], [16] which describe examples of compressed sensing for analog signals, and the work on finite-rate of innovation [17], [18]. In the last section of this article, we very briefly touch on this important area.

I. SAMPLING AND RECONSTRUCTION SETUP

The traditional Shannon sampling theorem (also attributed to Nyquist, Whittaker and Kotelnikov) states that a signal $x(t)$ bandlimited to π/T can be recovered from its uniform samples at time instances nT . The reconstruction is obtained by filtering with a sinc interpolation kernel:

$$x(t) = \frac{1}{T} \sum_{n=-\infty}^{\infty} x(nT) \text{sinc}(t/T - n),$$

where $\text{sinc}(t) = \sin(\pi t)/(\pi t)$. Although widely used, this theorem relies on three fundamental assumptions that are rarely met in practice. First, natural signals are almost never perfectly bandlimited. Second, the sampling device is usually not ideal, *i.e.*, it does not produce the exact values of the signal at the sampling locations. A common situation is that the ADC integrates the signal, usually over small neighborhoods around the sampling

TABLE I: Different scenarios treated in this review

	Unconstrained Reconstruction	Predefined Interpolation Kernel	Fine Grid Interpolation
Subspace Priors	Section II-A, II-B	Section II-C	Section II-D
Smoothness Priors	Section III-A	Section III-B	Section III-C
Stochastic Priors	Section IV-A	Section IV-B	Section IV-B

TABLE II: Design objective in each scenario

	Unconstrained Reconstruction	Predefined Interpolation Kernel	Fine Grid Interpolation
Subspace Priors	Perfect reconstruction	Minimum error	Minimum error
Smoothness Priors	Consistency/minimax	Consistency/minimax	Consistency/minimax
Stochastic Priors	MSE	MSE	MSE

points. Moreover, nonlinear distortions are often introduced during the sampling process. Finally, the use of the sinc kernel for reconstruction is often impractical due to its very slow decay.

To design interpolation methods that are adapted to practical scenarios, there are several issues that need to be properly addressed:

- 1) The sampling mechanism should be adequately modelled;
- 2) Relevant prior knowledge on the class of input signals should be accounted for;
- 3) Limitations should be imposed on the reconstruction algorithm in order to ensure robust and efficient recovery.

In this review we treat each of these three essential components of the sampling scheme. We focus on several models for each of the ingredients, which commonly arise in signal processing, image processing and communication systems. The setups we overview are summarized in Table I, and are detailed in the ensuing subsections. The design objective used in each scenario is indicated in Table II. As we discuss further below, the different priors dictate distinct objectives. For example, when the only information we have about the signal is that it is smooth, then the error cannot be minimized uniformly over all signals, and alternative design strategies are needed.

A. Sampling Process

1) *Linear Distortion*: In the Shannon sampling theorem, $x(t)$ is bandlimited to π/T and thus an equivalent strategy is to first filter the signal with a low-pass filter (LPF) with cut-off π/T and then uniformly sample the output. This interpretation is depicted in Fig. 1 with $s(-t) = \text{sinc}(t/T)$ being the impulse response of the LPF. The samples $c[n]$ can be expressed as

$$c[n] = \int_{t=-\infty}^{\infty} x(t)s(t-nT)dt \triangleq \langle x(t), s(t-nT) \rangle, \quad (1)$$

where $\langle y(t), h(t) \rangle$ denotes the $L_2(\mathbb{R})$ inner product between two finite-energy continuous-time signals. For simplicity, throughout the paper we assume a sampling interval of $T = 1$.




Fig. 1: A shift invariant sampling scheme. Filtering the signal $x(t)$ prior to taking ideal and uniform samples, can be restated as L_2 inner-products between the input $x(t)$ and $s(t)$. Shannon's paradigm corresponds to the special choice $s(-t) = \text{sinc}(t/T)$.




Fig. 2: A nonlinear and shift invariant sampling scheme. Prior to sampling with the functions $\{s(t - nT)\}$ the amplitudes of the signal $x(t)$ are distorted by the nonlinear mapping M .

In practical applications the sampling is not ideal. Therefore, a more appealing setting is to let $s(t)$ be an arbitrary sampling function. This allows to incorporate imperfections in the ideal sampler into the function $s(t)$ [19], [11], [20], [9]. As an example, typically the ADC averages the signal over a small interval rather than outputting point-wise signal values. This distortion can be taken into account by modifying $s(t)$ to include the integration.

2) *Nonlinear Distortion:* A more complicated situation arises when the sampling process includes nonlinear distortions. One simple approach to model nonlinearities is to assume that the signal is distorted by a memoryless, nonlinear invertible mapping prior to sampling by $s(-t)$, as in Fig. 2. This rather straightforward model is general enough to capture many systems of practical interest. Nonlinearities appear in a variety of setups and applications of digital signal processing including power electronics [21] and radiometric photography [22]. As another example, CCD image sensors introduce nonlinear distortions when excessive light intensity causes saturation. In some cases, nonlinearity is insinuated deliberately in order to increase the possible dynamic range of the signal while avoiding amplitude clipping, or damage to the ADC [23].

B. Signal Priors

In essence, the Shannon sampling theorem states that if $x(t)$ is known a priori to lie in the space of bandlimited signals, then it can be perfectly recovered from uniformly-spaced ideal samples. Clearly, the question of whether or not $x(t)$ can be recovered from its samples depends on the prior knowledge we have on the class of input signals. In this review we depart from the traditional bandlimited assumption and discuss signal priors that appear more frequently in signal processing and communication scenarios.

1) *Subspace Priors:* Our first focus is on signal spaces that are shift-invariant (SI). A SI subspace $\mathcal{A}$ of the space L_2 of bounded norm signals, with generator $a(t)$, is the space of signals that can be expressed as linear

combinations of shifts of $a(t)$:

$$x(t) = \sum_{n=-\infty}^{\infty} b[n]a(t-n), \quad (2)$$

where $b[n]$ is an arbitrary norm-bounded sequence. Note that $b[n]$ does not necessarily correspond to samples of the signal, *i.e.*, $x(n) \neq b[n]$ in general. Choosing $a(t) = \text{sinc}(t)$ results in the space of π -bandlimited signals. However, a much broader class of signal spaces can be defined including spline functions [8]. In these cases $a(t)$ can be easier to handle numerically than the sinc function.

A spline $f(t)$ of order N is a piecewise polynomial with the pieces combined at knots, such that the function is continuously differentiable $N-1$ times. It can be shown that any spline of order N with knots at the integers can be generated using (2) by a B -spline of order N which is the function obtained by the $(N+1)$ -fold convolution of the unit square

$$b(t) = \begin{cases} 1 & 0 < t < 1; \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Signals of the type (2) are also encountered when the analog signal to be sampled has originated from a digital source. For example, in communication systems, signals of this form are produced by pulse amplitude modulation transmitters. Extensive research in this field has been devoted to design receivers that undo the effect of inter-symbol interference, caused by the overlap of the pulses $a(t-n)$. Here we provide a geometric interpretation of this problem, which leads to insight into which classes of signals can be perfectly recovered from their samples. This viewpoint also allows to incorporate various constraints on the reconstruction method.

We note that although the discussion in this article is limited to SI subspaces, the results we present are valid in more general subspaces as well [9], [24].

2) *Smoothness Priors*: Subspace priors are very useful because, as we will see, they often can be utilized to perfectly recover $x(t)$ from its samples. However, in many practical scenarios our knowledge about the signal is much more limited and can only be formulated in very general terms. An assumption prevalent in image processing, is that natural signals are smooth in some sense. Here we focus on approaches that quantify the extent of smoothness using the L_2 norm $\|Lx(t)\|$, where L is usually chosen as some differential operator. The appeal of these models stems from the fact that they lead to linear recovery procedures. This is in contrast with smoothness measures such as total variation which result in nonlinear interpolation techniques.

The class of “smooth” signals is much richer than its subspace counterpart. Consequently, it is often impossible to distinguish between one “smooth” signal and another based solely on their samples. In other words, the sampling process inevitably entails information loss. Since perfect recovery cannot be attained in this scenario, we focus on two alternative criteria: consistency (or least-squares) and a worst case (minimax) design.

3) *Stochastic Priors*: The last family we consider in detail are stochastic priors. In this category, the signal $x(t)$ is modelled as a wide-sense stationary (WSS) random process with known power spectral density (PSD), a viewpoint prominent in the field of statistical signal processing. The design criterion here is the minimization of the mean-squared error (MSE) given the signal samples. The theory of sampling random signals has developed in parallel lines to its deterministic counterpart [5]. Interestingly, the stochastic setting leads to reconstruction




Fig. 3: Reconstruction using a digital compensation filter $h[n]$ and interpolation kernel $w(t)$.

techniques which are very similar to the methods arising from the smoothness priors. This provides an interesting equivalence between the smoothness operator L and the PSD of $x(t)$ in the random setup. Furthermore, we show that the study of statistical priors also sheds some light on the origin of artifacts, which are commonly encountered in traditional interpolation methods.

4) *Sparsity Priors*: In the last section we very briefly touch on sparsity priors. This class of functions lead to nonlinear reconstruction algorithms that have a very different structure than the linear interpolation methods in the majority of this paper. Since the treatment of these priors differs substantially from the rest of the review, we only point to several basic recovery techniques and results in this emerging area. A more detailed discussion merits a separate paper.

C. Reconstruction Methods

For a sampling theorem to be practical, it must take into account constraints that are imposed on the interpolation method. One aspect of the Shannon sampling theorem which renders it unrealizable, is the use of the sinc interpolation kernel. Due to its slow decay, the evaluation of $x(t)$ at a certain time instance t_0 , requires using a large number of samples located far away from t_0 . In many applications, reduction of computational load is achieved by employing much simpler methods, such as linear interpolation. In these cases the sampling scheme should be modified to compensate for the chosen non-ideal kernel.

1) *Unconstrained Reconstruction*: The first setup we consider is unconstrained recovery. Here, we design interpolation methods that are best adapted to the underlying signal prior according to the objectives in Table II, without restricting the reconstruction mechanism. In these scenarios, it is sometimes possible to obtain perfect recovery, as in the Shannon sampling theorem. The unconstrained reconstruction methods under the different scenarios treated in this paper (besides the case in which there are nonlinear distortions) all have a joint structure, depicted in Fig. 3. Here $w(t)$ is the impulse response of a continuous-time filter, which serves as the interpolation kernel, while $h[n]$ represents a discrete-time filter used to process the samples prior to reconstruction. Denoting the output of the discrete-time filter by $d[n]$, the input to the filter $w(t)$ is a modulated impulse train $\sum_n d[n]\delta(t - n)$ and the filter's output is given by

$$\hat{x}(t) = \sum_{n=-\infty}^{\infty} d[n]w(t - n). \quad (4)$$

The optimal interpolation kernels resulting from such considerations may be derived in the frequency domain but typically do not admit a closed form in the time domain. This limits the applicability of this approach to

situations where the kernel needs to be calculated only on a discrete set of points. In this case, the discrete Fourier transform (DFT) can be used to approximate the desired values. Consequently, apart from the case in which $w(t)$ is a spline, these methods seem to have been used *e.g.*, in the image processing community, only as a means of enlarging an image by an integer factor. More general geometrical transformations, such as rotation, lens distortion correction and scaling by an arbitrary factor, are typically not tackled using these methods.

2) *Predefined Kernel*: To overcome the difficulties in implementing the unconstrained solutions, we may resort to a system that uses a predefined interpolation kernel which is easy to implement. In this setup, the only freedom is in the design of the digital correction filter $h[n]$ in Fig. 3, which may be used to compensate for the non-ideal behavior of the pre-specified kernel $w(t)$ [19], [9], [12], [24], [11]. The filter $h[n]$ is selected to optimize a criterion matched to the signal prior.

By restricting the reconstruction to the form (4), we are essentially imposing that the recovered signal $\hat{x}(t)$ lie in the SI space generated by the pre-specified kernel $w(t)$. The class of SI spaces is very general and includes many signal spaces that lead to highly efficient interpolation methods. For example, by appropriate choice of $w(t)$ the family of splines can be described using (4). B-splines have been used for interpolation in the mathematical literature since the pioneering work of Schonberg [25]. In signal processing applications the use of B-splines gained popularity due to the work of Unser et. al. that showed how B-spline interpolation can be implemented efficiently [26], [27]. In particular, interpolation using splines of order up to 3 is very common in image processing. This is due to their tendency to efficiently represent smooth signals and the relatively low computational complexity needed for their evaluation at arbitrary locations.

3) *Fine Grid Interpolation*: The constraint to a predefined interpolation kernel may lead to severe degradation in the reconstruction. This emphasizes the fundamental trade-off between performance and implementation considerations. A common way to improve the performance of a reconstruction algorithm is to first upsample the digital signal and then apply some simple interpolation method on the resulting finer grid. This is a widely practiced approach for sampling rate conversion, where usually a rectangular or triangular interpolation filter is used [28].

Under mild conditions on the interpolation kernel, this approach allows to approximate the optimal unconstrained solution arbitrarily well by using a fine enough grid. This, of course, comes at the cost of computational complexity. In practice, it is not the asymptotic behavior that interests us, but rather optimizing the performance for a fixed setup. Thus, given a *fixed* over-sampling factor $K \geq 1$ and an interpolation filter $w(t)$, we would like to design a multirate system that processes the samples $c[n]$ and produces fine-grid expansion coefficients $d[n]$ such that the reconstruction

$$\hat{x}(t) = \sum_{n=-\infty}^{\infty} d[n] w\left(t - \frac{n}{K}\right) \quad (5)$$

well approximates $x(t)$. This setup is depicted in Fig. 4. Besides extending the discussion to general interpolation filters, we also relax the standard assumption that the input signal is bandlimited. Instead, we design a correction system that is best adapted to the prior we have on the input signal.

The interpolation methods corresponding to the different scenarios discussed above are summarized in Table III.




Fig. 4: Fine grid reconstruction using an upsampler followed by a digital compensation filter $h[n]$ and interpolation kernel $w(t)$. The rate of the sequence $d[n]$ is K times larger than that of $c[n]$.

TABLE III: Methods for signal recovery

	Unconstrained Reconstruction	Predefined Interpolation Kernel	Fine Grid Interpolation
Subspace Priors	Linear sampling: (12) Nonlinear distortion: (20), (21), (22)	(23)	(24)
Smoothness Priors	(28), (27)	Consistent: (32) Minimax: (34)	(38)
Stochastic Priors	(39), (40)	(41)	

II. SUBSPACE PRIORS

We begin by treating the setting in which the input signal $x(t)$ is known to lie in a given SI subspace. We show that when the reconstruction method is not restricted, these priors allow for perfect recovery of $x(t)$ from its nonideal samples both in the linear setting of Fig. 1 as well as in the presence of nonlinear distortions as in Fig. 2. Specifically, for any sampling function $s(t)$ there are a broad class of subspace priors under which $x(t)$ can be perfectly reconstructed. Conversely, for any given class of functions there are many choices of $s(t)$ that will allow for perfect recovery. These filters only have to satisfy a rather mild requirement. The surprising fact is that these results are valid even when a memoryless, invertible nonlinearity is inserted prior to sampling, as long as the nonlinearity does not vary too fast.

In the second part of this section, we extended the discussion to constrained reconstruction scenarios. In these cases perfect recovery is often impossible, as the restriction narrows down the set of candidate signals which the system can output. However, we will show that it is often possible to produce a reconstruction that minimizes the norm of the error.

Throughout this section $x(t)$ is assumed to lie in a subspace $\mathcal{A}$ generated by $a(t)$ (see (2)). In order for $\mathcal{A}$ to be well defined and the corresponding sampling theorems to be stable, the functions $\{a(t - n)\}$ should generate a Riesz basis or a frame [7]. To simplify the exposition we focus throughout on the case in which these functions are linearly independent and therefore form a basis. However, all the results extend easily to the case in which they are linearly dependent. In essence, a Riesz basis is a set of linearly independent vectors that ensures stable expansions, namely a small modification of the expansion coefficients results in a small modification of the signal (see Box C). In order for $a(t)$ to generate a Riesz basis the continuous-time Fourier transform (CTFT) of $a(t)$




Fig. 5: A unique element in $\mathcal{A}$ which is consistent with the samples in $\mathcal{S}$ can be recovered from the known samples.

must satisfy

$$\alpha \leq \sum_{k=-\infty}^{\infty} |A(\omega - 2\pi k)|^2 \leq \beta \quad \text{a.e. } \omega, \quad (6)$$

for some constants $\alpha > 0$ and $\beta < \infty$. The term in the middle of (6) is the discrete-time Fourier transform (DTFT) of the sampled correlation function $r_{aa}[n] = \langle a(t), a(t-n) \rangle$. More details on the CTFT and DTFT are given in Box A. In particular, the functions $\{a(t-n)\}$ form an orthonormal basis if and only if $\alpha = \beta = 1$ in (6).

A. Unconstrained Reconstruction with Linear Sampling

In the setup of Fig. 1 the input signal $x(t)$ is sampled by a set of sampling functions $\{s(t-n)\}$. We denote by $\mathcal{S}$ the space spanned by these sampling functions so that any $f(t)$ in $\mathcal{S}$ is of the form $f(t) = \sum_n d[n]s(t-n)$ for some bounded-norm sequence $d[n]$. We assume throughout that $s(t)$ satisfies the Riesz basis condition (6).

In order to understand what class of signals can be reconstructed from these samples we first observe that knowing the samples $c[n]$ is equivalent to knowing the orthogonal projection of $x(t)$ onto $\mathcal{S}$, which we denote by $x_{\mathcal{S}}(t) = P_{\mathcal{S}}x(t)$ (see Box B). Indeed,

$$c[n] = \langle x(t), s(t-n) \rangle = \langle x(t), P_{\mathcal{S}}s(t-n) \rangle = \langle P_{\mathcal{S}}x(t), s(t-n) \rangle, \quad (7)$$

where we used the fact that $P_{\mathcal{S}}s(t-n) = s(t-n)$. Since the functions $s(t-n)$ span $\mathcal{S}$, and $x_{\mathcal{S}}$ lies in $\mathcal{S}$, it is clear that $x_{\mathcal{S}}$ can be reconstructed from the samples $c[n]$. An immediate consequence is that if $x(t)$ lies in $\mathcal{S}$ so that $x(t) = x_{\mathcal{S}}(t)$, then it can be perfectly recovered.

This geometric interpretation implies that the question of reconstruction from $c[n]$ is equivalent to asking which signals can be recovered from knowledge of their orthogonal projection onto $\mathcal{S}$. At first glance it may seem like only signals in $\mathcal{S}$ may be reconstructed since the projection zeros out any component in $\mathcal{S}^\perp$. However, a closer inspection reveals that if we know in advance that $x(t)$ lies in a space $\mathcal{A}$ with suitable properties (which we will define below), then there is a unique vector in $\mathcal{A}$ with the given projection onto $\mathcal{S}$. As depicted in Fig. 5, in this case we can draw a vertical line from the projection until we hit the space $\mathcal{A}$ and in such a way obtain the unique vector in $\mathcal{A}$ that is consistent with the given samples. Thus perfect recovery is possible for a broad class of signals beyond those that lie in $\mathcal{S}$.

We next discuss how to recover $x(t)$ explicitly using a discrete-time filter as in Fig. 3. We first note that the orthogonal projection $P_S x(t)$ can be obtained from the samples $c[n]$ by using the scheme in Fig. 3 with $w(t) = s(t)$ and $h[n]$ chosen as the impulse response of the filter with DTFT [29], [8]

$$H(e^{j\omega}) = \frac{1}{\sum_{k=-\infty}^{\infty} |S(\omega - 2\pi k)|^2}, \quad (8)$$

where $S(\omega)$ is the CTFT of $s(t)$. Note that the Riesz basis condition (6) guarantees that (8) is well defined. Efficient implementation of (8), and the filters we introduce in the sequel, is possible in spline spaces, based on the results of [8], [26], [27].

To show that the output of the resulting system is $P_S x(t)$ note that if $x(t)$ is in $\mathcal{S}^\perp$, then the output will be zero since in this case $c[n]$ is the zero sequence. On the other hand, if $x(t) \in \mathcal{S}$, then we can write $x(t) = \sum_n b[n]s(t-n)$ for some sequence $b[n]$. Using the Fourier relations given in Box A it follows that

$$C(e^{j\omega}) = B(e^{j\omega}) \sum_{k=-\infty}^{\infty} |S(\omega - 2\pi k)|^2. \quad (9)$$

Therefore, $d[n] = b[n]$ and $\hat{x}(t) = x(t)$. Consequently, if $x(t)$ lies in $\mathcal{S}$ to begin with, then this scheme will ensure perfect reconstruction. If in addition $s(t)$ has the partition of unity property, i.e., $\sum_n s(t-n) = 1$ for all t , then it can be shown that by selecting the sampling period T sufficiently small, any input signal that is norm bounded can be approximated as close as desired by this scheme [8].

The denominator in (8) is the DTFT of the sampled correlation function $r_{ss}[n] = \langle s(t), s(t-n) \rangle$. Therefore, if the functions $\{s(t-n)\}$ form an orthonormal basis, then $r_{ss}[n] = \delta[n]$ and $H(e^{j\omega}) = 1$. In this case no pre-processing of the samples is necessary prior to reconstruction. This is precisely the setting of the Shannon sampling theorem: it is easy to verify that the functions $s(t-n) = \text{sinc}(t-n)$ form an orthonormal basis [?], [8].

To extend recovery beyond the space $\mathcal{S}$, suppose that $x(t)$ lies in a known subspace $\mathcal{A}$. Clearly in order to be able to reconstruct $x(t)$ from the given samples we need that $\mathcal{A}$ and $\mathcal{S}^\perp$ are disjoint since any signal $y(t)$ in the intersection of $\mathcal{A}$ and $\mathcal{S}^\perp$ will yield zero samples and therefore cannot be recovered. Intuitively, we also need $\mathcal{A}$ and $\mathcal{S}$ to have the same number of degrees of freedom. These requirements can be made precise by assuming a direct sum condition $L_2 = \mathcal{A} \oplus \mathcal{S}^\perp$ where $\oplus$ denotes a sum of two subspaces that intersect only at the zero vector. This implies that $\mathcal{A}$ and $\mathcal{S}^\perp$ are disjoint, and together span the space of L_2 signals. In the SI setting this condition translates into a simple requirement on the CTFT of the generators $a(t)$ and $s(t)$ [30]:

$$|\phi_{SA}(e^{j\omega})| > \alpha, \quad (10)$$

for some constant $\alpha > 0$, where

$$\phi_{SA}(e^{j\omega}) \triangleq \sum_{k=-\infty}^{\infty} S^*(\omega - 2\pi k) A(\omega - 2\pi k), \quad (11)$$

and $(\cdot)^*$ denotes the complex conjugate. The function $\phi_{SA}(e^{j\omega})$ is the DTFT of the sampled cross-correlation sequence $r_{sa} = \langle s(t), a(t-n) \rangle$ (See Box A). Under this condition, reconstruction can be obtained by choosing $w(t) = a(t)$ and [19], [12], [24], [11], [9]

$$H(e^{j\omega}) = \frac{1}{\phi_{SA}(e^{j\omega})}. \quad (12)$$

When $\mathcal{A} = \mathcal{S}$, the filter (12) coincides with (8).

To see that the filter (12) ensures perfect recovery for signals in $\mathcal{A}$, note that any $x(t) \in \mathcal{A}$ can be written as $x(t) = \sum_n b[n]a(t-n)$. Using the relations in Box A it can be shown that the sequence of samples will have a DTFT given by

$$C(e^{j\omega}) = B(e^{j\omega})\phi_{SA}(e^{j\omega}), \quad (13)$$

from which the result follows. In addition, for any $x(t) \in \mathcal{S}^\perp$ we have immediately that $\hat{x}(t) = 0$ since $c[n]$ will be the zero sequence. Consequently, the resulting system implements an oblique projection $E_{\mathcal{AS}^\perp}$ with range space $\mathcal{A}$ and null space $\mathcal{S}^\perp$ [31] (see Box B). Indeed this is the unique operator satisfying $E_{\mathcal{AS}^\perp}x(t) = x(t)$ for any $x(t)$ in $\mathcal{A}$, and $E_{\mathcal{AS}^\perp}x(t) = 0$ for $x(t)$ in $\mathcal{S}^\perp$.

It is also interesting to interpret the proposed sampling scheme as a basis expansion. Since any signal in $\mathcal{A}$ can be recovered from the corrected samples $d[n] = c[n] * h[n]$ via $x(t) = \sum_n d[n]a(t-n)$, we may view this sequence as the coefficients in a basis expansion. To obtain the corresponding basis we note that by combining the effects of the sampler $s(t)$ and the correction filter $h[n]$ of (12), the sequence of samples can be equivalently expressed as $d[n] = \langle x(t), v(t-n) \rangle$ where $v(t) = \sum_n h[n]s(t-n)$. In the Fourier domain,

$$V(\omega) = H(e^{j\omega})S(\omega). \quad (14)$$

Therefore, we conclude that any $x(t) \in \mathcal{A}$ can be written as

$$x(t) = \sum_{n=-\infty}^{\infty} \langle x(t), v(t-n) \rangle a(t-n). \quad (15)$$

It can be easily verified that the functions $\{v(t-n)\}$ form a Riesz basis for $\mathcal{S}$, and $\langle v(t-n), a(t-m) \rangle = \delta_{mn}$ where $\delta_{mn} = 1$ if $m = n$ and 0 otherwise. Therefore, these functions are the oblique dual basis of $\{a(t-n)\}$ in $\mathcal{S}$ [11], [12], [24], [30], [32], [33] (see Box C). When $\mathcal{A} = \mathcal{S}$, we recover the conventional dual basis functions. In this case $\{v(t-n)\}$ forms a basis for $\mathcal{S}$ that is dual to the original basis $\{s(t-n)\}$: $\langle v(t-n), s(t-m) \rangle = \delta_{nm}$. This provides a concrete method for constructing a dual of a given basis $\{a(t-n)\}$ in any subspace $\mathcal{S}$ satisfying the direct sum condition $L_2 = \mathcal{A} \oplus \mathcal{S}^\perp$.

To conclude our discussion so far, we have seen that a signal $x(t)$ in a SI subspace $\mathcal{A}$ generated by $a(t)$, can be reconstructed from its generalized samples in Fig. 1 using any choice of $s(t)$ for which (10) is satisfied. Thus for a given SI space, there is a broad variety of sampling filters we can select from. By choosing the functions appropriately, a variety of interesting sampling theories can be constructed, such as pointwise sampling of non-bandlimited signals, bandlimited sampling of nonbandlimited functions, and many more. An example is given below.

Despite the fact that any sampling function $s(t)$ satisfying (10) can be used to sample $x(t)$ in the space $\mathcal{A}$ generated by $a(t)$, in the presence of noise out of $\mathcal{A}$, the choice of sampling kernel will effect the reconstructed signal. More specifically, we have seen that the output of Fig. 3 with $w(t) = a(t)$ and $h[n]$ given by (12) is equal to the oblique projection $x_E(t) = E_{\mathcal{AS}^\perp}x(t)$. When $x(t) \in \mathcal{A}$, then $x_E(t) = x(t)$ for any choice of $\mathcal{S}^\perp$, or equivalently any sampling function $s(t)$ in Fig. 3 satisfying (10). However, if $x(t)$ does not lie entirely in

$\mathcal{A}$, for example due to noise, then different functions $s(t)$ will result in different approximations $x_E(t) \in \mathcal{A}$. A natural question is given an interpolation kernel $a(t)$, which choice of sampling function $s(t)$ will lead to a reconstruction $\hat{x}(t)$ that is closest to $x(t)$? If we measure the error using the squared-norm $\|\hat{x}(t) - x(t)\|^2$ where $\|f(t)\|^2 = \langle f(t), f(t) \rangle$, then the choice $s(t) = a(t)$ minimizes the error. This follows from the projection theorem which states that the orthogonal projection onto $\mathcal{A}$ is the closest vector in $\mathcal{A}$ to an arbitrary input $x(t)$ [34]:

$$\arg \min_{v(t) \in \mathcal{A}} \|x(t) - v(t)\|^2 = P_{\mathcal{A}}x(t). \quad (16)$$

Therefore, since $\hat{x}(t)$ is always in $\mathcal{A}$ the smallest error will result when $\hat{x}(t) = P_{\mathcal{A}}x(t)$ which can be achieved only if the sampling function $s(t)$ generates $\mathcal{A}$. In addition, in contrast to the orthogonal projection, an oblique projection can increase the norm of the noise at the input (see Box B). In practice however we may prefer other choices that are easier to implement at the expense of a slight increase in error [19], [9].

We conclude this subsection with a non-intuitive example in which a non-bandlimited signal is filtered with a LPF prior to sampling, and still can be perfectly reconstructed from the resulting samples. Suppose that the signal $x(t)$ is formed by exciting an RC circuit with a modulated impulse train $\sum_n d[n]\delta(t - n)$, as shown in Fig. 6. The impulse response of the RC circuit is known to be $a(t) = \tau^{-1} \exp\{-t/\tau\}u(t)$, where $u(t)$ is the unit step function and $\tau = RC$ is the time constant. Therefore

$$x(t) = \frac{1}{\tau} \sum_{n=-\infty}^{\infty} d[n] \exp\{-(t - n)/\tau\}u(t - n). \quad (17)$$

Clearly, $x(t)$ is not bandlimited. Now, suppose that $x(t)$ is filtered by an ideal low-pass filter $s(t) = \text{sinc}(t)$ and then sampled at times $t = n$ to obtain the sequence $c[n]$. Intuitively, there seems to be information loss in the sampling process since the entire frequency content of $x(t)$ outside $[-\pi, \pi]$ is zeroed out. This can be seen in the lower right-hand plot in Fig. 6. However, it is easily verified that condition (10) is satisfied in this setup and therefore perfect recovery is possible. The digital correction filter (12) in this case can be shown to be

$$h[n] = \begin{cases} 1 & n = 0; \\ \frac{\tau}{n}(-1)^n & n \neq 0. \end{cases} \quad (18)$$

Thus, to reconstruct $x(t)$ we need to excite an identical RC circuit with an impulse train modulated by the sequence $d[n] = h[n] * c[n]$.

B. Unconstrained Reconstruction with Nonlinear Distortion

Suppose now that as in the previous section $x(t)$ lies in a subspace $\mathcal{A}$ and (10) is satisfied. However, prior to sampling by $s(-t)$ the signal is distorted by a memoryless, nonlinear and invertible mapping $M(x)$ as in Fig. 2. A naive approach to recover the signal $x(t)$ from its samples is to first apply M^{-1} to the sample sequence $c[n]$ leading to a sequence $d[n]$, and then reconstruct $x(t)$ from the samples $d[n]$ using standard reconstruction techniques [35]. However, if the samples $c[n]$ are not ideal namely are not pointwise evaluations of $x(t)$, then this approach is suboptimal in general.

A surprising result developed in [36] is that if the nonlinearity is invertible and does not change too fast, then it does not introduce theoretical difficulties. More specifically, under the same direct sum condition (10) we had




Fig. 6: A non-bandlimited signal $x(t)$, formed by exciting an RC-circuit with a modulated impulse train, is sampled after passing through an ideal low-pass filter. Although a large portion of the frequency content is lost due to the filtering operation (bottom right), perfect recovery is still possible by digitally filtering the samples $c[n]$ and using the result to re-excite an identical RC-circuit.

in the linear sampling case, and bounds on the derivative of the nonlinearity, there is a unique signal $x(t)$ with the given samples $c[n]$. Therefore, it is enough to seek a recovery $\hat{x}(t)$ that is consistent in the sense that it yields the samples $c[n]$ after it is reinjected into the system: $\int_{-\infty}^{\infty} s(t - n)M(\hat{x}(t))dt = c[n]$. Any such signal must be equal to $x(t)$ due to the uniqueness property.

Since $x(t) \in \mathcal{A}$, we can write $x(t) = \sum_n d[n]a(t - n)$ for some sequence $d[n]$. Thus, our problem reduces to finding a sequence $d[n]$ which minimizes the consistency cost function

$$f(d) = \|c[n] - \hat{c}[n]\|_{\ell_2}, \quad (19)$$

with

$$\hat{c}[n] = \int_{-\infty}^{\infty} s(t - n)M\left(\sum_{m=-\infty}^{\infty} d[m]a(t - m)\right)dt, \quad (20)$$

where we know that the minimal value of $f(d)$ is 0. Since M is nonlinear, the cost function (19) is in general non-convex. Therefore optimization algorithms for minimizing (19) might trap a stationary point, and not the global minimum which we seek. Surprisingly, it can be shown [36] that under the direct sum condition and the derivative bounds on M , any stationary point of (19) is also a unique global minimum. Therefore, any algorithm designed to trap a stationary point automatically leads to perfect recovery. In Fig. 7 we show a block diagram




Fig. 7: One iteration of the nonlinear recovery algorithm. The expansion coefficients at the k th iteration, $d_k[n]$, are used to synthesize an estimate of the signal, $\hat{x}_k(t)$. This estimate goes through the sampling process to produce the corresponding samples $\hat{c}_k[n]$. The error with respect to the measured samples, $e_k[n] = c[n] - \hat{c}_k[n]$, is then used to update the estimate.

of an iterative approach which is derived by applying a Newton method on (19). This same algorithm can also be obtained from an approximate projection onto convex sets strategy, and a linearization approach; see [36] for more details.

At each iteration, the algorithm of Fig. 7 works as follows. Denote by $d_k[n]$ the expansion coefficients at the k th iteration. Then $d_{k+1}[n]$ is calculated as

$$d_{k+1}[n] = d_k[n] + \alpha_k \sum_{m=-\infty}^{\infty} g_k[n, m] e_k[m], \quad (21)$$

where α_k is the step size, $e_k[m] = c[m] - \hat{c}_k[m]$ is the error-in-samples sequence with $\hat{c}_k[n]$ denoting the approximate samples at stage k obtained via (20) with $d[n] = d_k[n]$, and $g_k[l, m]$ is a linear system which is the inverse of

$$h_k[l, m] = \int_{-\infty}^{\infty} s(t-l) M' \left(\sum_{n=-\infty}^{\infty} d_k[n] a(t-n) \right) a(t-m) dt. \quad (22)$$

Note that $h_k[l, m]$ is not shift invariant in general and therefore it cannot be inverted in the frequency domain to obtain $g_k[l, m]$. In practice, though, one usually analyzes a finite set of samples $c[n], 0 \leq n \leq N-1$. Assuming that $c[n] = 0$ outside this range, the matrix $\{g_k[l, m]\}$ for $0 \leq l, m \leq N-1$ can be obtained by inverting the corresponding matrix $\{h_k[l, m]\}$.

We now demonstrate the effectiveness of the algorithm in a scenario similar to that of Fig. 6. Specifically, suppose that, as in Fig. 6, $x(t)$ is known to be of the form (17), and we are given the samples $c[n] = \int_{n-1}^n \arctan(x(t)) dt$. Thus, the nonlinear mapping here is $M(x) = \arctan(x)$ and the sampling filter is a rectangular window of size 1. On the left hand side of Fig. 8, it can be seen that the samples $c[n]$ seem to constitute a rather poor representation of the signal. Consequently, if one ignores the nonlinearity and uses the techniques developed in the previous section, *i.e.*, filtering with $H(e^{j\omega})$ of (12), then the reconstruction error is large (dotted line). Another simple




Fig. 8: A signal $x(t)$ lying in a shift invariant space was linearly sampled after passing through a memoryless nonlinear system. Reconstruction techniques that ignore the nonlinear distortion or the linear sampler, fail to perfectly reconstruct $x(t)$ (left). However $x(t)$ can be perfectly recovered using the algorithm presented here (right).

approach is to first invert the nonlinearity and then filter the corrected samples $\tilde{c}[n] = \tan(c[n])$ with $H(e^{j\omega})$ of (12) (dashed line). Neither of these methods result in perfect recovery. On the right hand side of Fig. 8 we show the result of applying the algorithm presented here, which leads to perfect reconstruction of $x(t)$ from the nonideal samples $c[n]$.

C. Constrained Reconstruction

Up until now we specified the sampling process but did not restrict the reconstruction or interpolation kernel $w(t)$ in Fig. 3. In many applications this kernel is fixed in advance due to implementation issues. For example, in image processing applications kernels with small supports are often used. These include nearest neighbor, bilinear, bicubic, lanczos and splines. The interpolation kernel $w(t)$ can also represent the pixel shape of an image display.

In order to obtain stable reconstruction, we concentrate in the sequel on cases in which $w(t)$ satisfies the Riesz basis condition (6). In particular, it can be easily shown that B -splines all satisfy this requirement.

Given a sampling function $s(-t)$ and a fixed interpolation kernel $w(t)$ an important question is how to design the digital filter $h[n]$ in Fig. 3 so that the output $\hat{x}(t)$ is a good approximation of the input signal $x(t)$ in some sense. Clearly, $\hat{x}(t)$ will always lie in the space $\mathcal{W}$, spanned by the generator $w(t)$. This is because for every choice of the sequence $d[n]$, $\hat{x}(t)$ has the form $\hat{x}(t) = \sum_n d[n]w(t - n)$. Therefore, if $x(t)$ does not lie in $\mathcal{W}$ to begin with, then $\hat{x}(t)$ cannot be equal $x(t)$. Since $\hat{x}(t)$ is constrained to lie in $\mathcal{W}$, it follows from the projection theorem (16) that the minimal error approximation to $x(t)$ is obtained when $\hat{x}(t) = P_{\mathcal{W}}x(t)$. The question is whether this solution can be generated from the samples $c[n]$. In general, the answer is negative without sufficient prior knowledge on the signal [9]. However, when $x(t)$ lies in a subspace satisfying (10), $P_{\mathcal{W}}x(t)$ can be obtained

by filtering the sample sequence with

$$H(e^{j\omega}) = \frac{\phi_{WA}(e^{j\omega})}{\phi_{SA}(e^{j\omega})\phi_{WW}(e^{j\omega})}, \quad (23)$$

where $\phi_{WA}(e^{j\omega})$, $\phi_{SA}(e^{j\omega})$ and $\phi_{WW}(e^{j\omega})$ are as in (11) with the corresponding substitution of the filters $W(\omega)$, $A(\omega)$ and $S(\omega)$. In this case, the output of the system of Fig. 3 is given by $P_{\mathcal{W}}E_{\mathcal{AS}^\perp}x(t)$ [9]. Consequently, if $x(t) \in \mathcal{A}$, then $E_{\mathcal{AS}^\perp}x(t) = x(t)$ and the minimal squared-error approximation $P_{\mathcal{W}}x(t)$ is achieved.

To understand this result geometrically, note that we have already seen in the previous subsection that under the direct sum condition, any vector $x(t) \in \mathcal{A}$ can be recovered from the samples $c[n]$. This is illustrated in Fig. 5. Here, however, we are constrained to obtain a solution in $\mathcal{W}$. But, since we can determine $x(t)$, we can also compute $P_{\mathcal{W}}x(t)$, which is the minimal squared-error approximation in $\mathcal{W}$.

D. Dense Grid Recovery

The situation in which $x(t)$ can be completely determined from its samples but cannot be reproduced by the system is somewhat frustrating. Moreover, the error caused by restricting the recovered signal to lie in $\mathcal{W}$ may be very large if $\mathcal{W}$ is substantially different from $\mathcal{A}$. One way to bridge the gap between the unconstrained and constrained recovery techniques is to increase the interpolation rate, namely produce a reconstruction of the form $\hat{x}(t) = \sum_{n=-\infty}^{\infty} d[n]w(t - n/K)$, for some integer $K > 1$. This strategy is legitimate as we are still using a predefined interpolation kernel $w(t)$, which may be easy to implement. Thus, we effectively introduce a tradeoff between complexity and performance.

The motivation for this approach can be understood from a geometric viewpoint. As we increase the interpolation rate K , the reconstruction space $\mathcal{W}$ spanned by the functions $\{w(t - n/K)\}$ becomes “larger” and consequently “closer” to $\mathcal{A}$. In some cases, there exists a factor K for which $\mathcal{W}$ contains $\mathcal{A}$, thus recovering the possibility of perfect reconstruction.

In order for the reconstruction to be stable, we focus on the case in which the functions $\{w(t - n/K)\}$ form a Riesz basis. This requirement is satisfied if and only if there exists constants $0 < \alpha \leq \beta < \infty$ such that $\alpha \leq \sum_{l=-\infty}^{\infty} |W(\omega - 2\pi lK)|^2 \leq \beta$ is satisfied almost everywhere.

Similarly to the case in which $K = 1$, it can be shown that when $x(t)$ is in $\mathcal{A}$, the minimal squared error solution $\hat{x}(t) = P_{\mathcal{W}}x(t)$ can be attained with the system depicted in Fig. 4. The frequency response of the correction filter $h[n]$, which operates on the up-sampled data, is given by

$$H(e^{j\omega}) = \sum_{n=0}^{K-1} \sum_{m=0}^{K-1} \frac{\phi_{WA}(e^{j(K\omega+2\pi m)}) e^{\frac{2\pi j m n}{K}}}{\phi_{SA}(e^{jK\omega}) \phi_{WW}(e^{j(K\omega+2\pi m)})}, \quad (24)$$

where $\phi_{SA}(e^{j\omega})$ is defined in (11).

III. SMOOTHNESS PRIORS

Up until now we considered the setting in which the input signal $x(t)$ is constrained to a subspace. We now treat a more general and less restrictive formulation of the sampling problem in which our prior knowledge on the signal is that it is smooth in some sense. Here we model the extent of smoothness of $x(t)$ as the L_2 norm of the

output of the continuous-time filter $L(\omega)$ with $x(t)$ as its input. For example, the filter $L(\omega)$ can be chosen as a differential operator $L(\omega) = (1 + \omega^2)^\gamma$ for some choice of γ . We denote this quantity in operator form as $\|Lx\|$. For simplicity, we assume that $L(\omega) > \alpha > 0$ almost everywhere for some α , although the results extend to the non-invertible case as well.

Unlike subspace priors, a one-to-one correspondence between smooth signals and their sampled version does not exist since smoothness is a far less restrictive constraint than confining the signal to a subspace. This renders perfect recovery or even error norm minimization impossible. Indeed, it can be shown that there is no single choice of $\hat{x}(t)$ that minimizes $\|\hat{x}(t) - x(t)\|^2$ over all smooth signals $x(t)$. This is true even when $\hat{x}(t)$ is constrained to lie in a subspace $\mathcal{W}$. This is because the sample sequence $c[n]$ is no longer sufficient to determine the orthogonal projection $P_{\mathcal{W}}\hat{x}(t)$ [9]. Therefore, below, we focus on alternative approaches for designing the reconstruction system.

A. Unconstrained Reconstruction

To approximate $x(t)$ from its samples, based solely on the knowledge that it is smooth, we consider two design techniques. The first consists in finding the smoothest signal which gives rise to the measured samples $c[n]$ [19]. The second is a minimax strategy in which the system is designed to yield the best approximation for the worst-case signal among inputs that are consistent with the samples and are smooth to some extent [9].

1) *Smoothest Approximation:* In this approach we require that the reconstructed signal $\hat{x}(t)$ is smooth and consistent with the samples. The consistency requirement means that $\hat{x}(t)$ should yield the same samples $c[n]$ when reinjected into the system:

$$\langle \hat{x}(t), s(t - n) \rangle = c[n] = \langle x(t), s(t - n) \rangle \quad \text{for all } n. \quad (25)$$

The simplest strategy to produce a consistent smooth reconstruction is to minimize the smoothness $\|Lx\|$ subject to the consistency requirement:

$$\hat{x}(t) = \arg \min_{x(t)} \|Lx(t)\| \quad \text{subject to } S\{x(t)\} = c. \quad (26)$$

The notation $S\{x(t)\}$ denotes the sequence of samples $\langle s(t - n), x(t) \rangle$ and c stands for the sequence $\{c[n]\}$. It can be shown that the solution to (26) has the form of Fig. 3 where now the reconstruction kernel is

$$\tilde{W}(\omega) = \frac{S(\omega)}{|L(\omega)|^2}, \quad (27)$$

and

$$H(e^{j\omega}) = \frac{1}{\phi_{S\tilde{W}}(e^{j\omega})}. \quad (28)$$

As we have seen in the previous section, the filter (28) corresponds to the choice that leads to perfect reconstruction for signals $x(t) \in \tilde{\mathcal{W}}$ (see (8)). Thus, this approach can be viewed as first determining the optimal space given by (27), and then finding the unique signal in $\tilde{\mathcal{W}}$ that is consistent with the given samples.

As a special case, we may choose to produce the minimal norm consistent reconstruction $\hat{x}(t)$ by letting L be the identity operator I . This leads to $w(t) = s(t)$ and $H(e^{j\omega})$ is then given by (8). Consequently, $\hat{x}(t)$ is the orthogonal

projection onto the sampling space, $\hat{x}(t) = P_S x(t)$. This can also be seen by noting that any reconstruction $\hat{x}(t)$ which yields the samples $c[n]$ has the form $\hat{x}(t) = P_S x(t) + v(t)$ where $v(t)$ is an arbitrary vector in $\mathcal{S}^\perp$. Thus the minimal norm approximation corresponds to the choice $v(t) = 0$.

2) *Minimax Recovery*: The reconstruction error $\|\hat{x}(t) - x(t)\|^2$ of any recovery method depends on the unknown original signal $x(t)$. This renders the problem of comparing between interpolation methods complicated. Indeed, one algorithm may be better than another for certain input signals and worse for others. The next approach we discuss is based on optimizing the squared-error performance for the worst input signal.

The prior information we have can be used to construct a set $\mathcal{V}$ of all possible input signals:

$$\mathcal{V} = \{x(t) : S\{x(t)\} = c, \|Lx(t)\| \leq U\}, \quad (29)$$

where $U > 0$ is some finite constant. The set consists of signals that are consistent with the samples and are relatively smooth. We now seek the reconstruction that minimizes the worst-case error over $\mathcal{V}$:

$$\min_{\hat{x}(t)} \max_{x(t) \in \mathcal{V}} \|\hat{x}(t) - x(t)\|^2. \quad (30)$$

It can be shown that the optimal solution does not depend on the constant U . Furthermore, the minimax solution interestingly coincides with the smoothest approximation method, *i.e.*, the optimal interpolation kernel and digital compensation filter are given by (27) and (28) respectively.

Although the two approaches we discussed are equivalent in the unrestricted setting, the minimax strategy allows more flexibility in incorporating constraints on the reconstruction, as we show in the next subsection. Furthermore, it tends to outperform the consistency approach when further restrictions are imposed as we will demonstrate via several examples.

Figure 9 compares the minimax approach with bicubic interpolation in the context of image enlargement. The regularization operator was taken to be $L(\omega) = ((0.1\pi)^2 + \|\omega\|^2)^{1.3}$, where ω denotes the 2D frequency vector. In this example the minimax recovery is superior to the commonly used bicubic method in terms of peak signal to noise ratio (PSNR), defined as $\text{PSNR} = 10 \log_{10}(255^2/\text{MSE})$ with MSE denoting the empirical squared-error average over all pixel values. In terms of visual quality, the minimax reconstruction is sharper and contains enhanced textures.

B. Constrained Reconstruction

We next treat the problem of approximating $x(t)$ from its samples $c[n]$ using a pre-specified interpolation kernel $w(t)$. Similar to the unrestricted scenario, the two main approaches in this setup are consistent reconstruction [19], [8], [11], [12], [24] and minimax recovery [9], [37]. However, here the solutions no longer coincide. Both of these methods can be interpreted in terms of projections onto the spaces $\mathcal{W}$ and $\mathcal{S}$ that figure in the problem setting. Whereas the first approach leads to an oblique projection, the second strategy involves orthogonal projections which render it more robust to noise [38], [39].



Fig. 9: The Mandrill image was down-scaled by a factor of 3 using a rectangular sampling filter. The low resolution image was then enlarged back to its original dimensions using two interpolation methods. The bicubic interpolation kernel leads to a blurry reconstruction (left) with PSNR of 24.18dB. The minimax method leads to a sharper reconstruction (right) with PSNR of 24.39dB.

1) *Consistent Reconstruction*: In order to incorporate the constraint on the interpolation kernel, we extend (26) to include the restriction $x(t) \in \mathcal{W}$:

$$\hat{x}(t) = \arg \min_{x(t)} \|Lx(t)\| \text{ subject to } S\{x(t)\} = c, \quad x(t) \in \mathcal{W}. \quad (31)$$

Recall that under the direct sum condition (10) with $\mathcal{W}$ playing the role of $\mathcal{A}$, there is a unique signal in $\mathcal{W}$ satisfying $S\{x(t)\} = c$, which is equal to the oblique projection $E_{\mathcal{W}\mathcal{S}^\perp}x(t)$. Since there is only one signal in the constraint set of problem (31), the smoothness measure in the objective does not play a role. The oblique projection can be obtained by processing the samples $c[n]$ using the filter

$$H(e^{j\omega}) = \frac{1}{\phi_{SW}(e^{j\omega})}. \quad (32)$$

Comparing with (12), we see that this is precisely the filter that yields perfect recovery when we know that $x(t) \in \mathcal{W}$. When the direct sum condition is not satisfied, there can be several consistent solutions so that the objective in (31) is needed in order to select one output among all possibilities [40], [41].

2) *Minimax Recovery*: A drawback of the consistency approach is that the fact that $x(t)$ and $\hat{x}(t)$ yield the same samples does not necessarily imply that $\hat{x}(t)$ is close to $x(t)$. Indeed, for an input $x(t)$ not in $\mathcal{W}$, the norm of the resulting reconstruction error $\hat{x}(t) - x(t)$ can be made arbitrarily large, if $\mathcal{S}$ is close to $\mathcal{W}^\perp$. Furthermore, as we have seen, the consistency method essentially ignores the smoothness prior.

In order to directly control the reconstruction error $\|\hat{x}(t) - x(t)\|^2$, we suggest modifying the minimax strategy of the previous subsection to include the restriction $x(t) \in \mathcal{W}$. Therefore, our minimax design criterion is now:

$$\min_{\hat{x}(t) \in \mathcal{W}} \max_{x(t) \in \mathcal{V}} \|\hat{x}(t) - x(t)\|^2, \quad (33)$$

where $\mathcal{V}$ is the set of relatively-smooth consistent signals given by (29).

It turns out that the criterion (33) is too conservative and results in the trivial solution $\hat{x}(t) = 0$ [9]. To counterbalance the conservative behavior of the minimax approach, instead of minimizing the worst-case squared-norm error, we now consider minimizing the worst-case *regret* [42]. The regret is defined as the difference between the squared-norm error and the smallest possible error that can be achieved with a reconstruction in $\mathcal{W}$, namely $\|P_{\mathcal{W}^\perp}x(t)\|^2$. This error is attained when $\hat{x}(t) = P_{\mathcal{W}}x(t)$, which in general cannot be computed from the sequence of samples c . Since the regret depends in general on $x(t)$, it cannot be minimized for all $x(t)$. Instead we consider minimizing the worst-case regret over all possible values of $x(t)$ that are consistent with the given samples, which results in the problem

$$\min_{\hat{x}(t) \in \mathcal{W}} \max_{x(t) \in \mathcal{V}} \left\{ \|\hat{x}(t) - x(t)\|^2 - \|P_{\mathcal{W}^\perp}x(t)\|^2 \right\}, \quad (34)$$

with $\mathcal{V}$ given by (29). Expressing $x(t)$ as $x(t) = P_{\mathcal{W}}x(t) + P_{\mathcal{W}^\perp}x(t)$ and recalling that $\hat{x}(t) \in \mathcal{W}$ it is easy to see that (34) is equivalent to

$$\min_{\hat{x}(t) \in \mathcal{W}} \max_{x(t) \in \mathcal{V}} \|\hat{x}(t) - P_{\mathcal{W}}x(t)\|^2. \quad (35)$$

The solution to (35) can be shown to be the projection onto $\mathcal{W}$ of the unconstrained minimax recovery given by (27) and (28). This reconstruction can be obtained by digitally filtering the samples $c[n]$ with the filter

$$H(e^{j\omega}) = \frac{\phi_{W\tilde{W}}(e^{j\omega})}{\phi_{S\tilde{W}}(e^{j\omega}) \phi_{WW}(e^{j\omega})}, \quad (36)$$

where $\tilde{W}(\omega)$ is given by (27).

In Fig. 10 we demonstrate the difference between the consistent and minimax-regret methods in an image-enlargement task. The setup is the same as that of Fig. 9 only now the reconstruction filter is constrained to be a triangular kernel corresponding to linear interpolation. It can be seen that the error of the minimax regret recovery is only 0.7dB less than the unconstrained minimax shown in Fig. 9. The consistent approach, on the other hand, is much worse both in terms of PSNR and in terms of visual quality. Its tendency to over-enhance high frequencies stems from the fact that it ignores the smoothness prior.

Many of the interesting properties of the minimax-regret recovery can be best understood by examining the case where our only prior on the signal is that it is norm-bounded, *i.e.*, when L is the identity operator I . This scenario was thoroughly investigated in [9]. Setting $L(\omega) = 1$ in (28), the correction filter becomes

$$H(e^{j\omega}) = \frac{\phi_{SW}(e^{j\omega})}{\phi_{SS}(e^{j\omega}) \phi_{WW}(e^{j\omega})}. \quad (37)$$

Applying the Cauchy-Schwartz inequality to the numerator of (37) and to the denominator of (32), it is easy to see that the magnitude of the minimax regret filter (37) is smaller than that of the consistent filter (32) at all frequencies. This property renders the minimax regret approach more resistant to noise in the samples $c[n]$. This is because perturbations in $\hat{x}(t)$ caused by perturbations in $c[n]$ are always smaller in the minimax regret method than in the consistent approach.

Apart for robustness to digital noise, a disruption which takes place after the signal was sampled, the minimax regret method is also more resistant to perturbations in the continuous-time signal $x(t)$. To see this note that the



Fig. 10: The Mandrill image was down-scaled by a factor of 3 using a rectangular sampling filter. The low resolution image was then enlarged back to its original dimensions using the consistent and minimax regret methods. The consistent approach over-enhances the high frequencies (left) and results in a PSNR of 22.51dB. The minimax regret method leads to a smoother reconstruction (right) with PSNR of 23.69dB.

minimax regret reconstruction is given by $\hat{x}(t) = P_{\mathcal{W}}P_{\mathcal{S}}x(t)$. Thus, the norm of $\hat{x}(t)$ is necessarily bounded by that of $x(t)$. Furthermore, it is easy to show that the resulting reconstruction error is always bounded by twice the norm of $x(t)$: $\|\hat{x}(t) - x(t)\|^2 \leq 2\|x(t)\|^2$. In contrast, the consistent recovery is given by the oblique projection $\hat{x}(t) = E_{\mathcal{W}\mathcal{S}^\perp}x(t)$, which may increase the norm of $x(t)$. Consequently, the error of the consistent reconstruction can, in some cases, grow without bound.

In Fig. 11 we illustrate the minimax regret reconstruction geometrically for the case $L = I$. We have seen already that knowing the samples $c[n]$ is equivalent to knowing $x_{\mathcal{S}}(t) = P_{\mathcal{S}}x(t)$. In addition, our recovery is constrained to lie in the space $\mathcal{W}$. As illustrated in the figure, the minimax regret solution is a robust recovery scheme by which the signal is first orthogonally projected onto the sampling space, and then onto the reconstruction space.

Note that in the case in which $x(t)$ is known to lie in $\mathcal{S}$, it follows from the previous section that the minimal error can be obtained by using (23) with $A(\omega) = S(\omega)$. The resulting filter coincides with the minimax regret filter of (37). This implies that the regret approach minimizes the squared-error over all $x(t) \in \mathcal{S}$.

An interesting feature of the minimax regret solution is that it does not depend on the norm bound U . Therefore, $\hat{x}(t) = P_{\mathcal{W}}P_{\mathcal{S}}x(t)$ minimizes the worst-case regret error over all bounded inputs $x(t)$, regardless of the norm of $x(t)$. Furthermore, the regret recovery method does not require the direct-sum condition $L_2 = \mathcal{W} \oplus \mathcal{S}^\perp$, which is necessary in the development of the unique consistent approach.

In [9] tight bounds on the error resulting from each of the methods are developed and compared. We omit the technical details here and only summarize the main conclusions. We first recall that if we know a priori that $x(t)$ lies in a subspace $\mathcal{A}$ such that $L_2 = \mathcal{A} \oplus \mathcal{S}^\perp$, then the filter (23) will yield the minimal error approximation of $x(t)$ and therefore is optimal in the squared-norm sense. When $\mathcal{A} = \mathcal{S}$ this strategy reduces to the minimax regret




Fig. 11: Comparison of minimax regret reconstruction ($P_{\mathcal{W}}P_{\mathcal{S}}x(t)$) and consistent reconstruction ($E_{\mathcal{W}\mathcal{S}^+}x(t)$) for two different choices of $\mathcal{W}$. The minimax strategy is preferable when $\mathcal{W}$ is ‘far’ from $\mathcal{S}$ (left); Both methods lead to errors on the same order of magnitude when $\mathcal{W}$ is ‘close’ to $\mathcal{S}$ (right).

method, while if $\mathcal{A} = \mathcal{W}$, then we obtain the consistent reconstruction.

When no prior subspace knowledge is given, then the regret approach is preferable if the spaces $\mathcal{S}$ and $\mathcal{W}$ are sufficiently far apart, or if $x(t)$ has enough energy in $\mathcal{S}$. These results are intuitive as illustrated geometrically in Fig. 11. In Fig. 11(a) we depict the consistent and regret reconstruction when $\mathcal{W}$ is far from $\mathcal{S}$. As can be seen in the figure, in this case the error resulting from the consistent solution is large with respect to the regret approximation error. In Fig. 11(b), $\mathcal{W}$ and $\mathcal{S}$ are close, and the errors have roughly the same magnitude.

C. Minimax Dense Grid Reconstruction

We now extend the minimax regret approach to the dense-grid recovery setup of Fig. 4, in which the interpolation is performed using a predefined kernel $w(t)$ on a grid with n/K spacings. To treat this scenario within the minimax-regret framework, we need to solve problem (33) with the appropriate reconstruction space, namely $\mathcal{W} = \text{span}\{w(t - n/K)\}$. The corresponding correction filter can be shown to be

$$H(e^{j\omega}) = \sum_{n=0}^{K-1} \sum_{m=0}^{K-1} \frac{\phi_{W\tilde{W}}(e^{j(K\omega+2\pi m)}) e^{\frac{2\pi jmn}{K}}}{\phi_{S\tilde{W}}(e^{jK\omega}) \phi_{WW}(e^{j(K\omega+2\pi m)})}, \quad (38)$$

where $\tilde{W}(\omega)$ is given by (27).

To understand the necessity of fine grid interpolation, note that there is no analytic expression for the optimal unconstrained kernel (27) in the time domain. In rate conversion situations, where the output rate is an integer multiple of the input rate, the kernel $w(t)$ needs to be calculated only on a discrete set of points. This is because $\hat{x}(n/K) = \sum_m d[m]w(n/K - m)$, where K is the over-sampling factor. To approximate the sequence $\{w(n/K)\}$ on a finite set of indices, one can sample the expression $\sum_l W(\omega - 2\pi lK)$ on a finite set of frequencies and apply the inverse DFT. However, if $\hat{x}(t)$ needs to be evaluated at arbitrary locations, then this method cannot be used.

In the previous subsection we have seen that this problem can be tackled by using a predefined interpolation kernel for which a formula exists. An alternative approach is to first evaluate the optimal kernel (27) on a dense grid of instances and then use nearest neighbor or linear interpolation to obtain the values at the desired locations. This is called first and second order approximation [28]. Note, however, that this method does not take into account




Fig. 12: Comparison of first order approximation to the unconstrained optimal filter and dense grid minimax regret. The Mandrill image was enlarged by a factor of π/e in both methods. The reconstruction filter in both techniques is a triangular kernel corresponding to linear interpolation.

the non-optimal interpolation to be performed in the second stage. This is in contrast to the dense grid approach presented here. Filter (38) shapes the spectrum of the up-sampled sequence in a way that compensates for the non-optimal kernel to follow.

In Fig. 12 we compare the minimax-regret dense grid reconstruction approach and first order approximation to the unconstrained filter. To emphasize the differences, we used both methods to enlarge an image by a non-rational factor of π/e . It is clearly seen that the first-order approximation approach produces a blurry reconstruction whereas in the minimax method the edges are sharp and the textures are better preserved.

IV. STOCHASTIC PRIORS

In this section we treat signal priors of stochastic nature. Specifically, the input $x(t)$ is modelled as a WSS random process whose PSD is $\Lambda_{xx}(\omega)$. Our goal is to *linearly* estimate $x(t)$ given the samples $c[n]$. As we will see, the schemes resulting from these considerations have strong connections to the minimax methods of the previous section. In addition, this viewpoint also offers a nice explanation to reconstruction artifacts, frequently encountered in applications.

A. Unconstrained Reconstruction

We begin by examining the situation where there are no limitations on the reconstruction mechanism. In the deterministic setting with a smoothness prior we could not minimize the squared-error $\|\hat{x}(t) - x(t)\|^2$ for all smooth $x(t)$, and therefore instead discussed a minimax method. In contrast, in the stochastic setting we can use

the PSD of $x(t)$ in order to minimize the MSE $E[|x(t) - \hat{x}(t)|^2]$ for every t , which depends only on the statistics of $x(t)$ and not on the signal itself.

Our approach is to minimize the MSE by linearly processing the samples $c[n]$. As opposed to the common Wiener filtering problem, where both the input and output are either continuous- or discrete-time signals, here we are interested in estimating the continuous-time signal $x(t)$ based on equidistant samples of $y(t) = x(t) * s(-t)$. Consequently, we refer to this as the hybrid Wiener filtering problem.

It turns out that the reconstruction $\hat{x}(t)$ minimizing the MSE can be implemented by the block diagram in Fig. 3 with the interpolation kernel [43], [44], [45], [46]

$$W(\omega) = S(\omega)\Lambda_{xx}(\omega), \quad (39)$$

and digital correction filter

$$H(e^{j\omega}) = \frac{1}{\sum_{k=-\infty}^{\infty} |S(\omega - 2\pi k)|^2 \Lambda_{xx}(\omega - 2\pi k)}. \quad (40)$$

It is interesting to observe that (39) and (40) are identical to (27) and (28) with $\Lambda_{xx}(\omega) = |L(\omega)|^{-2}$. Therefore, the smoothness operator in the deterministic case corresponds to the whitening filter of the input $x(t)$ in the stochastic setting.

B. Constrained Reconstruction

We now treat the more practical constrained setting in which the interpolation filter is fixed in advance. Unfortunately, in this case, it turns out that for a general given interpolation kernel, there is no digital correction filter that can minimize the MSE for every t [47]. In fact, the filter minimizing the MSE at a certain time instance t_0 also minimizes the MSE at times $\{t_0 + n\}$ for all integer n , but not over the whole continuum. Therefore, error measures other than the pointwise MSE must be considered. Before treating the problem of choosing an appropriate criterion, we first discuss how this time dependence phenomenon is related to artifacts which are commonly encountered in certain interpolation methods.

The signal $x(t)$ in our setup is assumed to be WSS and, as a consequence, the sequence of samples $c[n]$ is a discrete WSS random process as is the output $d[n]$ of the digital correction filter in Fig. 4. The reconstruction $\hat{x}(t)$ is formed by modulating the shifts of the kernel $w(t)$ by the WSS discrete-time process $d[n]$. Assuming that the PSD of $d[n]$ is positive everywhere, signals of this type are *not stationary* unless $w(t)$ is π -bandlimited [48]. In practice, the interpolation kernels in use have a finite (and usually small) support, and thus are not bandlimited. In these cases, the periodic correlation in $\hat{x}(t)$ often degrades the quality of the reconstruction, as subjectively perceived by the visual or auditory system.

Note that despite the fact that natural signals are rarely stationary to begin with, it is still relevant to study how an interpolation algorithm reacts to stationary signals. In fact, if an interpolation scheme outputs a cyclostationary signal when fed with a stationary input, then it will commonly produce reconstructions with degraded subjective quality also when applied to real world signals, as demonstrated in Fig. 13. Note, however, that periodic structure in a recovered signal is not necessarily related to the MSE. For example, the optimal unrestricted kernel (39) is usually not bandlimited and thus leads to periodic structure in $\hat{x}(t)$.




Fig. 13: Periodic structure in an interpolated signal is a phenomena related to the effective bandwidth of the interpolation kernel. The larger the portion of its energy outside $[-\pi, \pi]$, the stronger the periodic correlation. The three images on the right were obtained by scaling a patch of the original image by a factor of 5 using three different methods. The portion of energy in the range $[-\pi, \pi]$ of the kernels is: rectangular - 61%, bicubic - 91%, sinc - 100%. Suppressed periodic correlation, however, does not necessarily imply that the reconstruction error is small.

The non stationary behavior of $\hat{x}(t)$ is the reason why the pointwise MSE can generally not be minimized for every t . Two alternative error measures that have been proposed are the sampling-period-average-MSE and the projected MSE.

The sampling-period-average-MSE utilizes the fact that the MSE is periodic, and integrates it over one period [49], [48]:

$$\text{MSE}_A = E \left[\int_{t_0}^{t_0+1} |x(t) - \hat{x}(t)|^2 dt \right]. \quad (41)$$

It turns out that the minimization of the average MSE leads to a correction filter independent of t_0 [48]. The second approach makes use of the fact that the best possible approximation to $x(t)$ in $\mathcal{W}$ is $P_{\mathcal{W}}x(t)$. Therefore this method aims at minimizing the projected MSE, defined as the MSE with respect to the optimal approximation in $\mathcal{W}$ [47]:

$$\text{MSE}_P = E [|P_{\mathcal{W}}x(t) - \hat{x}(t)|^2]. \quad (42)$$

Interestingly, both error measures (41) and (42) lead to the same digital correction filter, which is given by [48], [47]

$$H(e^{j\omega}) = \frac{\phi_{W\tilde{W}}(e^{j\omega})}{\phi_{S\tilde{W}}(e^{j\omega})\phi_{WW}(e^{j\omega})}, \quad (43)$$

where here $\tilde{W}(\omega) = S(\omega)\Lambda_{xx}(\omega)$. This is also the solution obtained by the minimax regret criterion (see (36)) where $|L(\omega)|^{-2}$ replaces the spectrum $\Lambda_{xx}(\omega)$. Therefore, here again, $L(\omega)$ plays the role of the whitening filter of $x(t)$.

The average MSE criterion (41) can also be used to handle the dense grid recovery setup of Fig. 4. The minimization of the average MSE, in this case, leads to the corresponding minimax regret solution with $\Lambda_{xx}(\omega) = |L(\omega)|^{-2}$.

V. SPARSITY PRIORS

Before concluding this review, we briefly address sampling of sparse signals, a topic that has gained considerable interest in recent years.

Throughout the review we focused mainly on linear interpolation techniques which were sufficient to recover, or properly approximate, many classes of signals. An important case where nonlinear methods are necessary is when the prior on the signal is not a subspace but rather a sparsity constraint. For example, we may know that the signal has the form

$$x(t) = \sum_{k=1}^{N-1} a_k g(t - t_k) \quad (44)$$

for some coefficients a_k and time instances t_k . Such a signal is said to have a sparse representation since only a few parameters are needed to specify it [17], [18]. If the values t_k are known then this is just a subspace prior. The more interesting scenario is when the times t_k need to be estimated along with the coefficients a_k . Several sampling strategies to deal with these signal classes have been suggested under the name of finite rate of innovation sampling. It turns out that roughly $2N$ samples are enough to recover the entire signal with proper post-processing.

Another important class of sparse signals are signals whose frequency transform (or any other transform) have a multiband structure. In this case the Fourier transform consists of a finite number N of bands, each of length at most B . If the band locations are known, then this corresponds to a subspace prior and the methods discussed in this review can be used to recover the signal from its samples. A very interesting question is whether the signal can be recovered at a rate lower than the Nyquist rate when the band locations are unknown. Such a sampling scheme is referred to as blind since it does not exploit knowledge of the band locations in the sampling and reconstruction stages.

At first site it may seem that since the band locations are unknown, the signal can have energy over the entire Nyquist frequency range, and therefore lower than Nyquist sampling will not be sufficient to recover the signal. Surprisingly, it turns out that this reasoning is in fact incorrect. In practice such classes of signals can be sampled at a rate much lower than the Nyquist rate, without impairing the ability to perfectly recover the signals [15]. The tradeoff is that the reconstruction involves nonlinear processing of the samples. In fact, it can be proven that the

minimal sampling rate at which such signals can be processed is twice the sum of the frequency bands, which is typically significantly smaller than twice the highest frequency, corresponding to the Nyquist rate. When the band locations are known, the minimal sampling rate is exactly the sum of the bands, referred to as the Landau rate. Thus, the price for constructing a blind system is only a factor of two.

The techniques developed to sample and reconstruct such classes of signals are based on ideas and results from the emerging field of compressed sensing [13], [14]. However, while the latter deals with sampling of finite vectors, the multiband problem is concerned with analog sampling. By using several tools, developed in more detail in [16], [50], it is possible to extend the essential ideas of compressed sensing to the analog domain. This broader framework combines ideas presented in this review with traditional compressed sensing tools in order to treat more general classes of signals, such as signals that lie in a union of shift-invariant spaces, *i.e.*,

$$x(t) = \sum_{k=1}^K \sum_{n=-\infty}^{\infty} d[n] a_k(t - n), \quad (45)$$

for a set of generators $a_k(t)$ where only $M < K$ out of the sequences $d[n]$ are not identically zero. This model subsumes the bandlimited class of functions as a special case.

APPENDIX A

BOX A: DISCRETE AND CONTINUOUS FOURIER TRANSFORMS

In this box we review the main properties of the discrete and continuous Fourier transforms used throughout the paper.

The continuous-time Fourier transform (CTFT) of a signal $x(t)$ in L_2 is defined as

$$X(\omega) = \int_{-\infty}^{\infty} x(t) e^{-j\omega t} dt. \quad (46)$$

We use the convention that upper case letters denote Fourier transforms. The discrete-time Fourier transform (DTFT) of a sequence $x[n]$ in ℓ_2 is defined by

$$X(e^{j\omega}) = \sum_{n=-\infty}^{\infty} x[n] e^{-j\omega n}. \quad (47)$$

The DTFT is 2π -periodic; to emphasize this fact we use the notation $X(e^{j\omega})$.

The DTFT of the sampled sequence $y(t = n)$ is related to the CTFT of $y(t)$ by

$$Y(e^{j\omega}) = \sum_{k=-\infty}^{\infty} Y(\omega - 2\pi k). \quad (48)$$

In the reverse direction, if the sequence $d[n]$ is used to create a continuous-time signal $f(t) = \sum_n d[n] y(t - n)$, then

$$F(\omega) = D(e^{j\omega}) Y(\omega). \quad (49)$$

An important sequence encountered in signal recovery problems is the sampled cross correlation $r_{as}[n] = \langle a(t), s(t - n) \rangle$. This sequence can be obtained by sampling the output of the filter $s(-t)$ with $a(t)$ as its input. An immediate consequence from (48) is that the DTFT of $r_{as}[n]$ can be expressed as

$$\phi_{SA}(e^{j\omega}) = \sum_{k=-\infty}^{\infty} S^*(\omega - 2\pi k) A(\omega - 2\pi k), \quad (50)$$




Fig. 14: Decomposition of a vector x into two components using an orthogonal projection (left) and an oblique projection (right).

where $(\cdot)^*$ denotes the complex conjugate. The set $\{a(t-n)\}$ is orthonormal if each function $a(t-n)$ is orthonormal to all of its integer shifts. This is equivalent to requiring that $r_{aa}[n] = \delta[n]$ where

$$\delta[n] = \begin{cases} 1 & n = 0; \\ 0 & \text{otherwise.} \end{cases} \quad (51)$$

From (50) we conclude that $\{a(t-n)\}$ is an orthonormal sequence if and only if

$$\phi_{aa}(e^{j\omega}) = \sum_{k=-\infty}^{\infty} |A(\omega - 2\pi k)|^2 = 1. \quad (52)$$

APPENDIX B

BOX B: PROJECTIONS IN HILBERT SPACES

A projection E in a Hilbert space $\mathcal{H}$ is a linear operator from $\mathcal{H}$ onto itself that satisfies the property

$$E^2 = E. \quad (53)$$

The importance of projection operators is that they map the entire space $\mathcal{H}$ onto the range space $\mathcal{R}(E)$, and leave vectors in this subspace unchanged. Furthermore, property (53) implies that every vector in $\mathcal{H}$ can be uniquely written as the combination of a vector in $\mathcal{R}(E)$ and a vector in the null space $\mathcal{N}(E)$, *i.e.*, we have the direct sum decomposition $\mathcal{H} = \mathcal{R}(E) \oplus \mathcal{N}(E)$. This is illustrated in Fig. 14 for two different projection operators. Therefore, a projection is completely determined by its range space and null space.

An orthogonal projection P is a Hermitian projection operator. In this case the range space $\mathcal{R}(P)$ and null space $\mathcal{N}(P)$ are orthogonal. Therefore, an orthogonal projection is completely determined by its range space. We use the notation $P_{\mathcal{V}}$ to denote an orthogonal projection with range $\mathcal{V} = \mathcal{R}(P_{\mathcal{V}})$.

An oblique projection $E_{\mathcal{V}\mathcal{W}}$ is an operator satisfying the projection property (53) that is not Hermitian. Its range space is given by $\mathcal{V}$ so that $E_{\mathcal{V}\mathcal{W}}x = x$ for any $x \in \mathcal{V}$, and its null space is given by $\mathcal{W}$ so that $E_{\mathcal{V}\mathcal{W}}x = 0$ for any $x \in \mathcal{W}$.

When decomposing the space using an orthogonal projection, the vectors comprising the decomposition are orthogonal, since $\mathcal{R}(P)$ and $\mathcal{N}(P)$ are orthogonal spaces. This is not true when using an oblique projection, as

illustrated in Fig. 14. Another important feature of the orthogonal projection is that the norm of the projection is never larger than the original norm:

$$\|P_V x\| \leq \|x\|. \quad (54)$$

This inequality does not necessarily hold for an oblique projection. In fact, the norm of the oblique projection can be much larger than the signal norm. Consequently, algorithms relying on the oblique projection can cause a significant increase in the noise if it is not constrained to the range space of the projection. On the other hand, orthogonal projections are more stable in the presence of noise due to (54).

APPENDIX C

BOX C: BASIS EXPANSIONS

A Schauder basis for a complex Hilbert space $\mathcal{H}$ is a countable set of vectors $\{x_n\}$ in $\mathcal{H}$ such that every vector $x \in \mathcal{H}$ can be written uniquely as a series

$$x = \sum_{n=-\infty}^{\infty} c[n]x_n \quad (55)$$

with scalars $c[n]$. For example, the set of complex exponentials $x_n(t) = \exp\{j\omega n t\}$ defined over $t \in [-\pi, \pi]$ is a Schauder basis for the space $L_2[-\pi, \pi]$ of square integrable functions over $[-\pi, \pi]$. In this basis, the expansion coefficients $c[n]$ of a function $x(t)$ are its Fourier coefficients.

A countable set of vectors $\{x_n\}$ in $\mathcal{H}$ is a Riesz basis for $\mathcal{H}$ if it is complete and there exist two constants $\alpha > 0$ and $\beta < \infty$ such that

$$\alpha \sum_n |c[n]|^2 \leq \left\| \sum_n c[n]x_n \right\|^2 \leq \beta \sum_n |c[n]|^2, \quad \forall c \in \ell_2. \quad (56)$$

Riesz bases have the desired stability property, namely that a slight change in the expansion coefficients $c[n]$ is ensured to entail only a small change in x . Consequently, these bases are important in ensuring stable sampling schemes.

An important question is how to obtain the expansion coefficients $c[n]$ of a vector x . If the basis $\{x_n\}$ is orthonormal (i.e., $\langle x_m, x_n \rangle = \delta_{mn}$) then $c[n] = \langle x, x_n \rangle$. This follows from taking the inner products of both sides of (55) with x_m and exploiting the orthogonality property. To determine the expansion coefficients when using a general non-orthogonal basis, we follow a similar route using the biorthogonal vectors, or dual basis $\tilde{x}_n$. The dual basis of x_n is the unique basis of $\mathcal{H}$ that satisfies the property

$$\langle x_m, \tilde{x}_n \rangle = \delta_{mn}. \quad (57)$$

Taking the inner products of both sides of (55) with respect to $\tilde{x}_m$, we find that

$$c[n] = \langle x, \tilde{x}_n \rangle. \quad (58)$$

If x_n is a Riesz basis, then so is its biorthogonal basis.

When the set of vectors $\{x_n\}$ span only a subspace $\mathcal{U}$ of $\mathcal{H}$, there may be many choices of biorthogonal bases in $\mathcal{H}$ satisfying (57). Intuitively, the biorthogonal basis vectors should span a subspace with the same number of

degrees of freedom as $\mathcal{U}$. A formal statement of this observation is that given any subspace $\mathcal{V}$ satisfying the direct sum condition $\mathcal{H} = \mathcal{U} \oplus \mathcal{V}^\perp$, there exists a unique set of vectors $\{\tilde{x}_n\}$ lying in $\mathcal{V}$ which constitute a biorthogonal basis for $\{x_n\}$. This set is called the oblique dual basis of $\{x_n\}$ in $\mathcal{V}$ [11], [12], [24], [30], [32], [33]. The vectors $\{\tilde{x}_n\}$ satisfy (57) and form a basis for $\mathcal{V}$, that satisfy the Riesz condition given that $\{x_n\}$ is a Riesz basis. In each subspace $\mathcal{V}$ there is only one dual basis. The canonical dual basis refers to the choice $\mathcal{U} = \mathcal{V}$. This concept can also be extended to sets of vectors that are linearly dependent, leading to oblique dual frames.

REFERENCES

- [1] T. Hentschel and G. Fettweis, "Sample rate conversion for software radio," *IEEE Communications Magazine*, vol. 38, no. 8, pp. 142–150, 2000.
- [2] T. Lehmann, C. Gonner, and K. Spitzer, "Survey: interpolation methods in medical image processing," *IEEE Trans. Medical Imaging*, vol. 18, no. 11, pp. 1049–1075, 1999.
- [3] C. E. Shannon, "Communications in the presence of noise," *Proc. IRE*, vol. 37, pp. 10–21, Jan 1949.
- [4] H. Nyquist, "Certain topics in telegraph transmission theory," *EE Trans.*, vol. 47, pp. 617–644, Jan. 1928.
- [5] A. J. Jerri, "The Shannon sampling theorem—Its various extensions and applications: A tutorial review," *Proc. of the IEEE*, vol. 65, no. 11, pp. 1565–1596, Nov. 1977.
- [6] S. G. Mallat, "A theory of multiresolution signal decomposition: The wavelet representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, pp. 674–693, 1989.
- [7] I. Daubechies, "The wavelet transform, time-frequency localization and signal analysis," *IEEE Trans. Inform. Theory*, vol. 36, pp. 961–1005, Sep. 1990.
- [8] M. Unser, "Sampling—50 years after Shannon," *IEEE Proc.*, vol. 88, pp. 569–587, Apr. 2000.
- [9] Y. C. Eldar and T. G. Dvorkind, "A minimum squared-error framework for generalized sampling," *IEEE Trans. Signal Processing*, vol. 54, no. 6, pp. 2155–2167, Jun. 2006.
- [10] T. Werther and Y. C. Eldar, "General framework for consistent sampling in Hilbert spaces," *International Journal of Wavelets, Multiresolution, and Information Processing*, vol. 3, no. 3, pp. 347–359, Sep. 2005.
- [11] Y. C. Eldar, "Sampling and reconstruction in arbitrary spaces and oblique dual frame vectors," *J. Fourier Anal. Appl.*, vol. 1, no. 9, pp. 77–96, Jan. 2003.
- [12] —, "Sampling without input constraints: Consistent reconstruction in arbitrary spaces," in *Sampling, Wavelets and Tomography*, A. I. Zayed and J. J. Benedetto, Eds. Boston, MA: Birkhäuser, 2004, pp. 33–60.
- [13] D. L. Donoho, "Compressed sensing," *IEEE Trans. on Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr 2006.
- [14] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [15] M. Mishali and Y. C. Eldar, "Blind multi-band signal reconstruction: Compressed sensing for analog signals," *CCIT Report no. 639, EE Dept., Technion - Israel Institute of Technology*; submitted to *IEEE Trans. Signal Process.*, Sep. 2007.
- [16] —, "Reduce and boost: Recovering arbitrary sets of jointly sparse vectors," *CCIT Report no. 686, EE Dept., Technion - Israel Institute of Technology*; submitted to *IEEE Trans. Signal Process.*, Feb. 2008.
- [17] M. V. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Processing*, vol. 50, pp. 1417–1428, June 2002.
- [18] P. Dragotti, M. Vetterli, and T. Blu, "Sampling moments and reconstructing signals of finite rate of innovation: Shannon meets Strang-Fix," *IEEE Trans. Sig. Proc.*, vol. 55, no. 5, pp. 1741–1757, May 2007.
- [19] M. Unser and A. Aldroubi, "A general sampling theory for nonideal acquisition devices," *IEEE Trans. Signal Processing*, vol. 42, no. 11, pp. 2915–2925, Nov. 1994.
- [20] P. P. Vaidyanathan, "Generalizations of the sampling theorem: Seven decades after Nyquist," *IEEE Trans. Circuit Syst. I*, vol. 48, no. 9, pp. 1094–1109, Sep. 2001.

- [21] F. Pejović and D. Maksimović, "A new algorithm for simulation of power electronic systems using piecewise-linear device models," *IEEE Tran. power elec.*, vol. 10, no. 3, pp. 340–348, 1995.
- [22] H. Farid, "Blind inverse gamma correction," *IEEE Tran. Image Proc.*, vol. 10, pp. 1428 – 1433, Oct 2001.
- [23] K. Kose, K. Endoh, and T. Inouye, "Nonlinear amplitude compression in magnetic resonance imaging: Quantization noise reduction and data memory saving," *IEEE AES magazine*, pp. 27–30, Jun 1990.
- [24] Y. C. Eldar and T. Werther, "General framework for consistent sampling in Hilbert spaces," *International Journal of Wavelets, Multiresolution, and Information Processing*, vol. 3, no. 3, pp. 347–359, Sep. 2005.
- [25] I. J. Schoenberg, *Cardinal Spline Interpolation*. Philadelphia, PA: SIAM, 1973.
- [26] M. Unser, A. Aldroubi, and M. Eden, "B-Spline signal processing: Part I - Theory," *IEEE Trans. Signal Processing*, vol. 41, no. 2, pp. 821–833, Feb 1993.
- [27] —, "B-Spline signal processing: Part II - Efficient design and applications," *IEEE Trans. Signal Processing*, vol. 41, no. 2, pp. 834–848, Feb 1993.
- [28] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms, and Applications*. Prentice-Hall, Inc. Upper Saddle River, NJ, USA, 1996.
- [29] I. Djokovic and P. P. Vaidyanathan, "Generalized sampling theorems in multiresolution subspaces," *IEEE Trans. Signal Processing*, vol. 45, pp. 583–599, Mar. 1997.
- [30] O. Christansen and Y. C. Eldar, "Oblique dual frames and shift-invariant spaces," *Applied and Computational Harmonic Analysis*, vol. 17, no. 1, 2004.
- [31] W. S. Tang, "Oblique projections, biorthogonal Riesz bases and multiwavelets in Hilbert space," *Proc. Amer. Math. Soc.*, vol. 128, no. 2, pp. 463–473, 2000.
- [32] Y. C. Eldar and O. Christansen, "Characterization of oblique dual frame pairs," *J. Applied Signal Processing*, pp. 1–11, 2006, article ID 92674.
- [33] O. Christensen and Y. C. Eldar, "Generalized shift-invariant systems and frames for subspaces," *J. Fourier Analys. Appl.*, vol. 11, pp. 299–313, 2005.
- [34] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont MA: Athena Scientific, 1999.
- [35] Y. M. Zhu, "Generalized sampling theorem," *IEEE Tran. Circ. Sys. II*, vol. 39, pp. 587–588, Aug 1992.
- [36] T. G. Dvorkind, Y. C. Eldar, and E. Matusiak, "Nonlinear and non-ideal sampling: Theory and methods," submitted to *IEEE Trans. Signal Processing*.
- [37] T. G. Dvorkind, H. Kirshner, Y. C. Eldar, and M. Porat, "Minimax approximation of representation coefficients from generalized samples," *IEEE Trans. Signal Processing*, vol. 55, pp. 4430–4443, Sep. 2007.
- [38] Y. Eldar and M. Unser, "Nonideal sampling and interpolation from noisy observations in shift-invariant spaces," *IEEE Trans. Signal Processing*, vol. 54, no. 7, pp. 2636–2651, Jul. 2006.
- [39] Y. C. Eldar, "Mean-squared error sampling and reconstruction in the presence of noise," *IEEE Trans. Signal Process*, vol. 54, no. 12, pp. 4619–4633, Dec. 2006.
- [40] A. Hirabayashi and M. Unser, "Consistent sampling and signal recovery," *IEEE Trans. on Sig. Proc.*, vol. 55, no. 8, pp. 4104–4115, August 2007.
- [41] T. G. Dvorkind and Y. C. Eldar, "Robust and consistent solutions for generalized sampling," submitted to *IEEE Signal Proc. Lett.*
- [42] Y. C. Eldar, A. Ben-Tal, and A. Nemirovski, "Linear minimax regret estimation of deterministic parameters with bounded data uncertainties," *IEEE Trans. Signal Processing*, vol. 52, pp. 2177–2188, Aug. 2004.
- [43] F. O. Huck, C. L. Fales, N. Halyo, R. W. Samms, and K. Stacy, "Image gathering and processing- Information and fidelity," *Optical Society of America, Journal, A: Optics and Image Science*, vol. 2, pp. 1644–1666, 1985.
- [44] M. B. Matthews, "On the linear minimum-mean-squared-error estimation of an undersampled wide-sense stationary random process," *IEEE Trans. Signal Processing*, vol. 48, no. 1, pp. 272–275, 2000.
- [45] C. A. Glasbey, "Optimal linear interpolation of images with known point spread function," *Scandinavian Image Analysis Conference—SCIA-2001*, pp. 161–168, 2001.

- [46] S. Ramani, D. Van De Ville, and M. Unser, "Sampling in practice: Is the best reconstruction space bandlimited?" in *Proceedings of the 2005 IEEE International Conference on Image Processing (ICIP'05)*, vol. II, Genova, Italy, September 11-14, 2005, pp. 153–156.
- [47] Y. C. Eldar and M. Unser, "Nonideal sampling and interpolation from noisy observations in shift-invariant spaces," *IEEE Trans. Signal Processing*, vol. 54, no. 7, pp. 2636–2651, 2006.
- [48] T. Michaeli and Y. C. Eldar, "High Rate Interpolation of Random Signals from Nonideal Samples," *Accepted to IEEE Trans. Signal Processing*, 2008.
- [49] T. Blu and M. Unser, "Quantitative Fourier analysis of approximation techniques: Part I - Interpolators and projectors," *IEEE Trans. Signal Processing*, vol. 47, no. 10, pp. 2783–2795, Oct. 1999.
- [50] Y. C. Eldar, "Compressed sensing of analog signals," in preperation.