Convex Optimization in Signal Processing and Communications

March 16, 2009

Contents

List of contributors

Tomer Michaeli
Tomer Michaeli received the B.Sc. degree in Electrical Engineering in 2004 (*Summa Cum Laude*) from the Technion–Israel Institute of Technology, Haifa, Israel. He is currently pursuing the Ph.D. degree in Electrical Engineering at the Technion. From 2000 to 2008, he was a research engineer at RAFAEL Research Laboratories, Israel Ministry of Defense, Haifa. In 2008 he held the Andrew and Erna Finci Viterbi Fellowship. His research interests include statistical signal processing, estimation theory, sampling theory, image processing and computer vision.

Yonina C. Eldar
Yonina C. Eldar received the B.Sc. degree in Physics in 1995 and the B.Sc. degree in Electrical Engineering in 1996 both from Tel-Aviv University (TAU), Tel-Aviv, Israel, and the Ph.D. degree in Electrical Engineering and Computer Science in 2001 from the Massachusetts Institute of Technology (MIT), Cambridge. From January 2002 to July 2002 she was a Postdoctoral Fellow at the Digital Signal Processing Group at MIT. She is currently an Associate Professor in the Department of Electrical Engineering at the Technion - Israel Institute of Technology, Haifa, Israel. She is also a Research Affiliate with the Research Laboratory of Electronics at MIT. Her research interests are in the general areas of sampling theory, statistical signal processing, and computational biology. Dr. Eldar was in the program for outstanding students at TAU from 1992 to 1996. In 1998, she held the Rosenblith Fellowship for study in Electrical Engineering at MIT, and in 2000, she held an IBM Research Fellowship. From 2002-2005 she was a Horev Fellow of the Leaders in Science and Technology program at the Technion and an Alon Fellow. In 2004, she was awarded the Wolf Foundation Krill Prize for Excellence in Scientific Research, in 2005 the Andre and Bella Meyer Lectureship, in 2007 the Henry Taub Prize for Excellence in Research, and in 2008 the Hershel Rich Innovation Award, the Award for Women with Distinguished Contributions, and the Muriel & David Jacknow Award for Excellence in Teaching. She is a member of the IEEE Signal Processing Theory and Methods technical committee and the Bio Imaging Signal Processing technical committee, an Associate Editor for the IEEE Transactions on Signal Processing, the EURASIP Journal of Signal Processing, the SIAM Journal on Matrix Anal-

ysis and Applications, and the SIAM Journal on Imaging Sciences, and on the Editorial Board of Foundations and Trends in Signal Processing.

Part I

1 Optimization Techniques in Modern Sampling Theory

Tomer Michaeli and Yonina C. Eldar
Department of Electrical Engineering, Technion–Israel Institute of Technology, Haifa, Israel

Sampling theory has benefited from a surge of research in recent years, due in part to intense research in wavelet theory and the connections made between the two fields. In this chapter we present several extensions of the Shannon theorem, which treat a wide class of input signals as well as nonideal sampling and constrained recovery procedures. This framework is based on an optimization viewpoint, which takes into account both the goodness of fit of the reconstructed signal to the given samples, as well as relevant prior knowledge on the original signal. Our exposition is based on a Hilbert-space interpretation of sampling techniques, and relies on the concepts of bases (frames) and projections. The reconstruction algorithms developed in this chapter lead to improvement over standard interpolation approaches in signal and image processing applications.

1.1 Introduction

Sampling theory treats the recovery of a continuous-time signal from a discrete set of measurements. This field attracted significant attention in the engineering community ever since the pioneering work of Shannon [1] (also attributed to Whitaker [2], Kotelnikov [3], and Nyquist [4]) on sampling bandlimited signals. Discrete-time signal processing (DSP) inherently relies on sampling a continuous-time signal to obtain a discrete-time representation. Therefore, with the rapid development of digital applications, the theory of sampling has gained importance.

Traditionally, sampling theories addressed the problem of perfectly reconstructing a given class of signals from their samples. During the last two decades, it has been recognized that these theories can be viewed in a broader sense of projection onto appropriate subspaces of L_2 [5, 6, 7], and also extended to arbitrary Hilbert spaces [8, 9].

The goal of this chapter is to introduce a complementary viewpoint on sampling, which is based on optimization theory. The key idea in this approach is to construct an optimization problem that takes into account both the goodness of fit of the reconstructed signal to the given samples, as well as relevant prior knowledge on the original signal, such as smoothness. This framework is rooted in the theory of spline interpolation: one of the arguments in favor of

using smoothing splines for interpolation is that they minimize an energy functional. This objective accounts for both the miss-fit at the sampling locations, and for the energy of the mth derivative [10] (where m is related to the spline order). Several other interpolation techniques have been proposed in recent years, which are based on variational arguments of the same spirit [11]. These methods also have connections with Wiener's estimation theory of random processes [11, 12, 13, 14, 15]. This chapter provides extensions, generalizations, and rigorous proofs of a variety of optimization-based interpolation techniques. Some of these methods were recently reported (without proof) in the review paper [15].

We focus on sampling problems in an abstract Hilbert space setting. This facilitates the treatment of various sampling scenarios via an optimization framework. Since signals are generally functions over a continuous domain, the optimization problems we will encounter in this chapter are infinite-dimensional and cannot be solved numerically. Furthermore, in order to make the discussion relevant to arbitrary Hilbert spaces and not only for signals which are functions over $\mathbb{R}$, we refrain from using calculus of variations in our derivations. Therefore, most of our treatment relies on the interpretation of the resulting optimization problems in terms of projections onto appropriate spaces. Besides being infinite-dimensional, many of the optimization problems we attempt to solve are also not convex. Nevertheless, we derive closed-form solutions for many of these, by employing a geometric viewpoint. In the last section we tackle interpolation problems which do not admit a closed form solution. In this scenario, we narrow the discussion to signals lying in $\mathbb{R}^n$ or $\mathbb{C}^n$ and employ semi-definite relaxation and saddle-point techniques to arrive at optimization problems which can be solved numerically.

The scenarios treated in this chapter differ from one another in several aspects. First, we distinguish between noiseless and noisy samples. Second, the reconstruction algorithms we consider can either be adjusted according to some objective, or constrained to be of a certain predefined structure. Third, we treat two types of prior knowledge on the original signal, which we term subspace priors and smoothness priors. Last, we treat two classes of optimization criteria for each of the scenarios: least-squares and minimax. The setups we consider are summarized in Table 1.1. Throughout the chapter we highlight the connection between the resulting reconstruction methods, demonstrate how they can be implemented efficiently, and provide concrete examples of interpolation results.

The chapter is organized as follows. In Section 1.2 we provide mathematical preliminaries needed for the derivations to follow. Section 1.3 describes in detail the sampling and reconstruction setups treated in this chapter. In particular, we elaborate on the types of prior knowledge and reconstruction approaches that are considered. Section 1.4 is devoted to the different objectives that are at the heart of the proposed recovery techniques. In Sections 1.5 and 1.6 we develop reconstruction methods for the case of subspace and smoothness priors respectively. Each of these priors is studied in a constrained and unconstrained reconstruction setting using both the least-squares and minimax objectives. All reconstruction methods in these sections possess closed-form expressions. Section 1.7 includes

Table 1.1: Different scenarios treated in this chapter

	Unconstrained Reconstruction		Constrained Reconstruction	
	Least-Squares	Minimax	Least-Squares	Minimax
Subspace Priors Noise-Free Samples	Section 1.5.1.1	Section 1.5.1.2	Section 1.5.2.1	Section 1.5.2.2
Smoothness Priors Noise-Free Samples	Section 1.6.1.1	Section 1.6.1.2	Section 1.6.2.1	Section 1.6.2.2
Smoothness Priors Noisy Samples	Section 1.8	Section 1.8	Section 1.8	Section 1.8

comparisons between the various recovery algorithms. Finally, in Section 1.8 we treat the case in which the samples are noisy. There, we focus our attention on smoothness priors and on the minimax objective. We use semi-definite relaxation to tackle the resulting non-convex quadratic programs. This section also includes a summary of recent results on semi-definite relaxation of non-convex quadratic programs, which is needed for our derivations.

1.2 Notation and mathematical preliminaries

The exposition in this chapter is based on a Hilbert-space interpretation of sampling techniques, and relies on the concepts of frames and projections. In this section we introduce some notations and mathematical preliminaries, which form the basis for the derivations in the sections to follow.

1.2.1 Notation

We denote vectors in an arbitrary Hilbert space $\mathcal{H}$ by lowercase letters, and the elements of a sequence $c \in \ell_2$ by $c[n]$. Traditional sampling theories deal with signals, which are defined over the real line. In this case the Hilbert space $\mathcal{H}$ of signals of interest is the space L_2 of square integrable functions, *i.e.,* every vector $x \in \mathcal{H}$ is a function $x(t)$, $t \in \mathbb{R}$. We use the notations x and $x(t)$ interchangeably according to the context. The continuous-time Fourier transform (CTFT) of a signal $x(t)$ is denoted by $X(\omega)$ and is defined by

$$X(\omega) = \int_{-\infty}^{\infty} x(t) e^{j\omega t} dt. \tag{1.1}$$

Similarly, the discrete-time Fourier transform (DTFT) of a sequence $c[n]$ is denoted by $C(e^{j\omega})$ and is defined as

$$C(e^{j\omega}) = \sum_{n=-\infty}^{\infty} c[n] e^{j\omega n}. \tag{1.2}$$

The inner product between vectors $x, y \in \mathcal{H}$ is denoted $\langle x, y \rangle$, and is linear in the second argument; $\|x\|^2 = \langle x, x \rangle$ is the squared norm of x. The ℓ_2 norm

of a sequence is $\|c\|^2 = \sum_n |c[n]|^2$. The orthogonal complement of a subspace $\mathcal{A}$ is denoted by $\mathcal{A}^\perp$. A direct sum between two closed subspaces $\mathcal{W}$ and $\mathcal{S}$ is written as $\mathcal{W} \oplus \mathcal{S}$, and is the sum set $\{w + s;\ w \in \mathcal{W}, s \in \mathcal{S}\}$ with the property $\mathcal{W} \cap \mathcal{S} = \{0\}$. Given an operator T, T^* is its adjoint, and $\mathcal{N}(T)$ and $\mathcal{R}(T)$ are its null space and range space, respectively.

1.2.2 Projections

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

$$E^2 = E. \tag{1.3}$$

A projection operator maps the entire space $\mathcal{H}$ onto the range $\mathcal{R}(E)$, and leaves vectors in this subspace unchanged. Property (1.3) implies that every vector in $\mathcal{H}$ can be uniquely decomposed into a vector in $\mathcal{R}(E)$ and a vector in $\mathcal{N}(E)$, that is, we have the direct sum decomposition $\mathcal{H} = \mathcal{R}(E) \oplus \mathcal{N}(E)$. 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, and consequently P is completely determined by its range. We use the notation $P_\mathcal{V}$ to denote an orthogonal projection with range $\mathcal{V} = \mathcal{R}(P_\mathcal{V})$. An important property of an orthogonal projection onto a closed subspace $\mathcal{V}$ is that it maps every vector in $\mathcal{H}$ to the vector in $\mathcal{V}$ which is closest to it:

$$P_\mathcal{V} y = \arg\min_{x \in \mathcal{V}} \|y - x\|. \tag{1.4}$$

This property will be useful in a variety of different sampling scenarios.

An oblique projection is a projection operator that is not necessarily Hermitian. The notation $E_{\mathcal{A}\mathcal{S}^\perp}$ denotes an oblique projection with range space $\mathcal{A}$ and null space $\mathcal{S}^\perp$. If $\mathcal{A} = \mathcal{S}$, then $E_{\mathcal{A}\mathcal{S}^\perp} = P_\mathcal{A}$ [16]. The oblique projection onto $\mathcal{A}$ along $\mathcal{S}^\perp$ is the unique operator satisfying

$$\begin{aligned} &E_{\mathcal{A}\mathcal{S}^\perp} a = a \text{ for any } a \in \mathcal{A}; \\ &E_{\mathcal{A}\mathcal{S}^\perp} s = 0 \text{ for any } s \in \mathcal{S}^\perp. \end{aligned} \tag{1.5}$$

Projections can be used to characterize the pseudo-inverse of a given operator. Specifically, let T be a bounded operator with closed range. The *Moore-Penrose pseudo inverse* of T, denoted $T^\dagger$, is the unique operator satisfying [17]:

$$\begin{aligned} &\mathcal{N}(T^\dagger) = \mathcal{R}(T)^\perp, \\ &\mathcal{R}(T^\dagger) = \mathcal{N}(T)^\perp, \\ &TT^\dagger x = x \quad \forall x \in \mathcal{R}(T). \end{aligned} \tag{1.6}$$

The following is a set of properties of the pseudo-inverse operator, which will be used extensively throughout the chapter [17].

Lemma 1.1. *Let T be a bounded operator with closed range. Then:*

1. $P_{\mathcal{R}(T)} = T^\dagger T$.
2. $P_{\mathcal{N}(T)^\perp} = TT^\dagger$.
3. T^* *has closed range and* $(T^*)^\dagger = (T^\dagger)^*$.

1.2.3 Frames

As we will see in the sequel, sampling can be viewed as the process of taking inner products of a signal x with a sequence of vectors $\{a_n\}$. To simplify the derivations associated with such sequences, we use the notion of set transforms.

Definition 1.1. *A set transformation $A : \ell_2 \to \mathcal{H}$ corresponding to vectors $\{a_n\}$ is defined by $Ab = \sum_n b[n]a_n$ for all $b \in \ell_2$. From the definition of the adjoint, if $c = A^*x$, then $c[n] = \langle a_n, x\rangle$.*

Note that a set transform A corresponding to vectors $\{a_n\}_{n=1}^N$ in $\mathbb{R}^M$ is simply an $M \times N$ matrix whose columns are $\{a_n\}_{n=1}^N$.

To guarantee stability of the sampling theorems we develop, we concentrate on vector sets that generate frames [18, 19].

Definition 1.2. *A family of vectors $\{a_n\}$ in a Hilbert space $\mathcal{H}$ is called a frame for a subspace $\mathcal{A} \subseteq \mathcal{H}$ if there exist constants $\alpha > 0$ and $\beta < \infty$ such that the associated set transform A satisfies*

$$\alpha\|x\|^2 \le \|A^*x\|^2 \le \beta\|x\|^2, \quad \forall x \in \mathcal{A}. \tag{1.7}$$

The norm in the middle term is the ℓ_2 norm of sequences.

The lower bound in (1.7) ensures that the vectors $\{a_n\}$ span $\mathcal{A}$. Therefore, the number of frame elements, which we denote by N, must be at least as large as the dimension of $\mathcal{A}$. If $N < \infty$, then the right hand inequality of (1.7) is always satisfied with $\beta = \sum_n \|a_n\|^2$. Consequently, any finite set of vectors that spans $\mathcal{A}$ is a frame for $\mathcal{A}$. For an infinite set of frame vectors $\{a_n\}$, condition (1.7) ensures that the sum $x = \sum_n b[n]a_n$ converges for any sequence $b \in \ell_2$ and that a small change in the expansion coefficients b results in a small change in x [20]. Similarly, a slight perturbation of x will entail only a small change in the inner products with the frame elements.

1.3 Sampling and reconstruction setup

We are now ready to introduce the sampling and reconstruction setup that will be studied in this chapter. Before we elaborate on the abstract Hilbert space

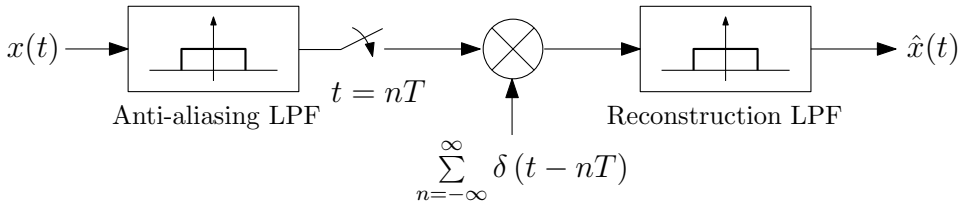

Figure 1.1: Shannon's sampling paradigm. The signal $x(t)$ passes through an ideal LPF prior to sampling. Reconstruction is obtained using the same LPF.

exposition that will be at the heart of our derivations to follow, we first review the famous Shannon sampling theorem [1].

Shannon's theorem states that a signal $x(t)$ bandlimited to π/T can be recovered from its uniform samples at time instants nT. Reconstruction is obtained by filtering the samples with a sinc interpolation kernel:

$$\hat{x}(t) = \frac{1}{T} \sum_{n=-\infty}^{\infty} x(nT)\text{sinc}(t/T - n), \tag{1.8}$$

where $\text{sinc}(t) = \sin(\pi t)/(\pi t)$. If the bandwidth of $x(t)$ exceeds π/T, then an anti-aliasing low-pass filter (LPF) with cutoff π/T can be used prior to sampling, as shown in Fig. 1.1. In this case, the reconstruction formula (1.8) produces the best approximation (in an L_2 sense) to $x(t)$ within the class of π/T-bandlimited functions. To see this, note that a LPF with unit magnitude in its pass-band satisfies (1.3) and is thus a projection operator. Furthermore, the range space of such a filter comprises all signals whose CTFT vanishes outside $[-\pi/T, \pi/T]$, and its null space is the set of signals that vanish in $[-\pi/T, \pi/T]$, which is orthogonal to the range. Therefore, the anti-aliasing filter is an *orthogonal* projection. Property (1.4) then implies that its output is the best approximation to its input within all signals in its range, namely π/T-bandlimited signals.

Shannon's theorem contains four fundamental aspects that are important in any sampling theorem:

1. **Prior knowledge:** The input lies in the class of π/T-bandlimited signals;
2. **Sampling mechanism:** Pre-filtering with a LPF with cutoff π/T, followed by pointwise sampling;
3. **Reconstruction method:** Sinc interpolation kernel modulated by the sample values;
4. **Objective:** Minimization of the L_2 norm of the error $x(t) - \hat{x}(t)$.

These specific choices of prior knowledge, sampling mechanism, reconstruction method, and objective are often not met in practical scenarios. First, natural signals are rarely truly bandlimited. Second, the sampling device is usually not ideal, that is, it does not produce exact signal values at the sampling locations. A common situation is that the analog-to-digital converter (ADC) integrates

the signal, usually over small neighborhoods surrounding the sampling points. Moreover, in many applications the samples are contaminated by noise due to quantization and other physical effects. Third, the use of the sinc kernel for reconstruction is often impractical because of its slow decay. Finally, when considering signal priors which are richer than the bandlimited assumption, it is usually impossible to minimize the error norm uniformly over all feasible signals. Therefore other criteria must be considered.

In this chapter we treat each of these essential components of the sampling scheme, focusing on several models which commonly arise in signal processing, image processing and communication systems. For simplicity, throughout the chapter we assume a sampling period of $T = 1$. We next elaborate on the signal priors and general sampling and reconstruction processes we treat.

1.3.1 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 ideal uniformly-spaced samples. Clearly, the question of whether $x(t)$ can be recovered from its samples depends on the prior knowledge we have on the class of input signals. In this chapter we depart from the traditional bandlimited assumption and discuss signal priors that appear more frequently in signal processing and communication scenarios.

1.3.1.1 Subspace priors

Our first focus is on cases where the signal $x(t)$ is known to lie in a subspace $\mathcal{A}$, spanned by vectors $\{a_n\}$. Although the discussion in this chapter is valid for a wide class of such subspaces, we take special interest in subspaces of L_2 that are shift invariant (SI). A SI subspace $\mathcal{A}$ of L_2, is a space of signals that can be expressed as linear combinations of shifts of a generator $a(t)$ [7]:

$$x(t) = \sum_{n=-\infty}^{\infty} d[n]a(t-n), \tag{1.9}$$

where $d[n]$ is an arbitrary norm-bounded sequence. Note that $d[n]$ does not necessarily correspond to samples of the signal, that is we can have $x(n) \neq d[n]$. More generally, $\mathcal{A}$ may be generated by several generators $a_k(t)$ so that $x(t) = \sum_{k=1}^{K} \sum_{n=-\infty}^{\infty} d_k[n]a_k(t-n)$. For simplicity we focus here on the single generator case. Using set transform formulation, (1.9) can be written compactly as

$$x = Ad, \tag{1.10}$$

where A is the set transform associated with the functions $\{a(t-n)\}$.

Choosing $a(t) = \mathrm{sinc}(t)$ in (1.9) results in the space of π-bandlimited signals. However, a much broader class of signal spaces can be defined including spline

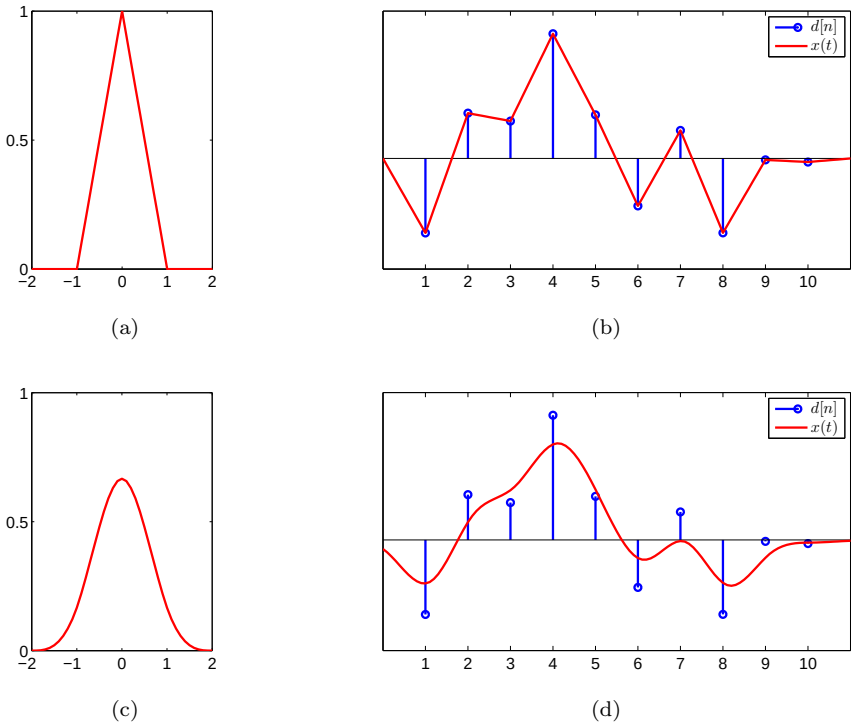

Figure 1.2: Spline functions of different orders generated using (1.9) with the same sequence $d[n]$. (a) $\beta_1(t)$. (b) A spline of degree 1. (c) $\beta_3(t)$. (d) A spline of degree 3.

functions [21]. In these cases $a(t)$ may be easier to handle numerically than the sinc function.

A popular choice of SI spaces in many image processing applications is the class of splines . A spline $f(t)$ of degree 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 degree N with knots at the integers can be generated using (1.9) by a B-spline of degree N, denoted $\beta_N(t)$. The latter is the function obtained by the $(N+1)$-fold convolution of the unit square

$$\beta_0(t) = \begin{cases} 1 & -\frac{1}{2} < t < \frac{1}{2}; \\ 0 & \text{otherwise.} \end{cases} \tag{1.11}$$

As demonstrated in Fig. 1.2, the sequence $d[n]$ in (1.9) is not equal to the samples $x(n)$ for splines of order greater than 1.

An important generalization of the SI subspace prior (1.9) is the class of signals that lie in a *union* of SI spaces. In this case,

$$x(t) = \sum_{k=1}^{K} \sum_{n=-\infty}^{\infty} d_k[n] a_k(t-n), \tag{1.12}$$

for a set of generators $a_k(t)$ where only $M < K$ out of the sequences $d_k[n]$ are not identically zero. However, we do not know in advance which M are chosen. This model can be used, for example, to describe multiband signals whose total number of active bands is small compared to the Nyquist rate [22, 23]. The techniques developed to sample and reconstruct such classes of signals are based on ideas and results from the emerging field of compressed sensing [24, 25]. 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 [26, 23, 27], it is possible to extend the essential ideas of compressed sensing to the analog domain. These results can also be applied more generally to signals that lie in a union of subspaces [28, 29], which are not necessarily shift invariant. Unlike subspace priors, nonlinear techniques are required in order to recover signals of this type. Therefore, for simplicity we will confine the discussion in this chapter to the single subspace case.

1.3.1.2 Smoothness priors

Subspace priors are very useful because, as we will see, they often can be used to perfectly recover $x(t)$ from its samples. However, in many practical scenarios our knowledge about the signal is much less complete and can only be formulated in very general terms. An assumption prevalent in image and signal 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 $\|Lx\|$, where L is a linear operator. Specifically, we assume that

$$\|Lx\| \leq \rho \tag{1.13}$$

for some finite constant $\rho > 0$. It is common to use SI smoothness priors, resulting in a linear time-invariant (LTI) operator L. In thess cases the corresponding filter $L(\omega)$ is often chosen to be a first or second order derivative in order to constrain the solution to be smooth and nonoscillating, *i.e.,* $L(\omega) = a_0 + a_1 j\omega + a_2(j\omega)^2 + \cdots$ for some constants a_n. Another common choice is the filter $L(\omega) = (a_0^2 + \omega^2)^\gamma$ with some parameter γ. The latter is in use mainly in image processing applications. The appeal of the model (1.13) stems from the fact that it leads to linear recovery procedures. In contrast, smoothness measures such as total variation [30], 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, in contrast to subspace priors, perfect reconstruction is typically impossible under the smooth-

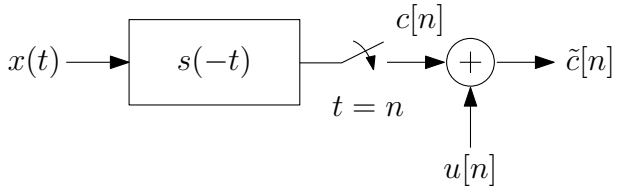

Figure 1.3: Shift-invariant sampling. Filtering the signal $x(t)$ prior to taking ideal and uniform samples, can be interpreted as L_2 inner-products between $x(t)$ and shifts of $s(t)$. In practical applications the samples are contaminated by noise $u[n]$.

ness assumption. Instead, we develop recovery algorithms that attempt to best approximate a smooth signal from the given samples.

1.3.2 Sampling process

We now present the general sampling process we treat. As we have seen, in the Shannon sampling theorem $x(t)$ is filtered with an LPF with cutoff π prior to sampling. In practical applications the sampling is not ideal. Therefore, a more realistic setting is to let the anti-aliasing filter, which we denote by $s(-t)$, be an arbitrary sampling function, as depicted in Fig. 1.3. This allows to incorporate imperfections in the ideal sampler into the function $s(t)$ [5, 8, 31, 32, 15]. As an example, typical ADCs average the signal over a small interval rather than outputting pointwise signal values. This distortion can be taken into account by modifying $s(t)$ to include the integration.

The samples $c[n]$ can be expressed as

$$c[n] = \int_{-\infty}^{\infty} x(t)s(t-n)dt = \langle s_n, x \rangle, \tag{1.14}$$

where $s_n(t) = s(t-n)$. More generally, we treat the scenario in which the samples $c[n]$ are obtained as inner products with a set of arbitrary functions $\{s_n\}$. Using set transform notation, the samples can be written as

$$c = S^* x, \tag{1.15}$$

where S is the set transform corresponding to the sampling vectors $\{s_n\}$ and c is the sequence whose nth element is $c[n]$.

To ensure that the sampling process is stable, we concentrate on the case in which the vectors $\{s_n\}$ form a frame for their span, which we term the sampling space $\mathcal{S}$. It follows immediately from the upper bound in (1.7) that the sequence of samples $c[n] = \langle s_n, x \rangle$ is then in ℓ_2 for any signal x that has bounded norm.

In the case where $\mathcal{S}$ is a SI space, condition (1.7) can be stated in terms of $S(\omega)$, the CTFT of the generator $s(t)$. Specifically, it can be shown that the

functions $\{s(t-n)\}$ generate a frame if and only if

$$\alpha \leq \phi_{SS}\left(e^{j\omega}\right) \leq \beta, \quad \omega \in \mathcal{I}_{\mathcal{S}}, \tag{1.16}$$

for some constants $\alpha > 0$ and $\beta < \infty$. Here,

$$\phi_{SS}\left(e^{j\omega}\right) = \sum_{k=-\infty}^{\infty} |S(\omega - 2\pi k)|^2 \tag{1.17}$$

is the DTFT of the sampled correlation function $r_{ss}[n] = \langle s(t), s(t-n)\rangle$ and $\mathcal{I}_{\mathcal{S}}$ is the set of frequencies ω for which $\phi_{SS}(e^{j\omega}) \neq 0$ [33]. It easy to see that $s(t) = \mathrm{sinc}(t)$ satisfies (1.16). Furthermore, B-splines of all orders also satisfy (1.16) [21].

In many situations the samples are perturbed by the sampling device, for example due to quantization or noise. Thus, as shown in Fig. 1.3, one usually only has access to the modified samples

$$\tilde{c}[n] = c[n] + u[n], \tag{1.18}$$

where $u[n]$ is a discrete-time noise process. Clearly, the noise should be taken into consideration when designing a reconstruction algorithm.

Another setup which was treated recently is that of reconstructing a signal which has undergone nonlinear distortion prior to sampling [34]. Using optimization tools together with frame-perturbation theory it can be shown that under a subspace prior and several technical conditions the signal can be reconstructed perfectly despite the nonlinearity. In this chapter, we focus on linear sampling and therefore do not survey these results.

1.3.3 Reconstruction method

The problem at the heart of sampling theory is how to reconstruct a signal from a given set of samples. For a sampling theorem to be practical, it must take into account constraints that are imposed on the reconstruction 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 instant 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 nearest-neighbor interpolation. In these cases the sampling scheme should be modified to compensate for the chosen non-ideal kernel.

In this chapter we study two interpolation strategies: unconstrained and constrained reconstruction. In the former, we pose no limitation on the interpolation algorithm. The goal then is to extend Shannon's theorem to more general classes of input signals. In the latter strategy, we restrict the reconstruction to be of a predefined form in order to reduce computational load. Here too, we address a variety of input signals.

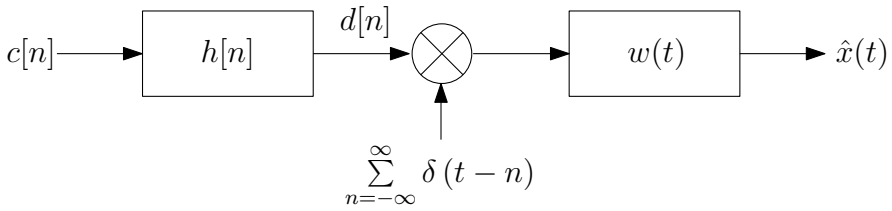

Figure 1.4: Reconstruction using a digital compensation filter $h[n]$ and interpolation kernel $w(t)$.

1.3.3.1 Unconstrained reconstruction

The first setup we treat is unconstrained recovery. Here, we design reconstruction methods that are best adapted to the underlying signal prior according to an appropriately defined objective, without restricting the reconstruction mechanism . In these scenarios, it is sometimes possible to obtain perfect recovery, as in the Shannon sampling theorem. When both the sampling process and signal prior are SI, the unconstrained reconstruction methods under the different scenarios treated in this chapter all have a common structure, depicted in Fig. 1.4. 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 analog filter $w(t)$ is a modulated impulse train $\sum_n d[n]\delta(t-n)$. The filter's output is given by

$$\hat{x}(t) = \sum_{n=-\infty}^{\infty} d[n]w(t-n). \tag{1.19}$$

If either the sampling process or the prior are not SI, then the reconstruction has the more general form $\hat{x} = \sum_n d[n]w_n$, where $d[n]$ is obtained by a linear transform of $c[n]$. Using set transform notation, the recovered signal can be written as

$$\hat{x} = Wd = WHc, \tag{1.20}$$

where W is the set transform corresponding to the reconstruction vectors $\{w_n\}$ and H is a linear operator. Our goal is to determine both the reconstruction vectors $\{w_n\}$ and the transform H to be optimal in some sense.

We will see that in the SI case, optimal interpolation kernels are typically derived in the frequency domain but very often do not admit a closed-form in the time domain. This limits the applicability of these recovery techniques. One way to resolve this problem is to choose the signal prior so as to yield an efficient interpolation algorithm, as done *e.g.,* in [12] in the context of exponential B-splines. Nevertheless, this approach restricts the type of priors that can be handled.

1.3.3.2 Constrained reconstruction

To overcome the difficulties in implementing the unconstrained solutions, we may resort to a system that uses a predefined interpolation kernel $w(t)$ that is easy to implement. For example, in image processing applications kernels with small supports are often used. These include nearest neighbor (B-spline of order 0), bilinear (B-spline of order 1), bicubic (Keys kernel [35]), Lanczos and B-spline of order 3. In this setup, the only freedom is in the design of the digital correction filter $h[n]$ in Fig. 1.4, which may be used to compensate for the non-ideal behavior of the pre-specified kernel [5, 32, 9, 36, 8]. The filter $h[n]$ is selected to optimize a criterion matched to the signal prior.

As a simple example demonstrating the need for constraining the reconstruction, consider the situation where we are given pointwise samples of a π-bandlimited signal and our goal is to compute $x(t_0)$ for some non-integer t_0. The sinc interpolation formula (1.8) uses the entire sequence of samples $c[n]$, since $\mathrm{sinc}(t)$ is not compactly supported. To reduce computational load, we can replace the sinc kernel, for instance, by $\beta_3(t)$ (a B-spline of degree 3). Since the support of the latter is $[-2, 2]$, the approximation (1.19) includes only 4 summands per time instance t_0. In this example, however, using the sample values $c[n]$ as the expansion coefficients $d[n]$ in (1.19) is not desired, as demonstrated in Fig. 1.2(d). To obtain a good approximation, a digital processing step is required prior to reconstruction, as depicted in Fig. 1.4.

More generally, when the signal is reconstructed using an arbitrary given set of vectors $\{w_n\}$, the goal is to design the correction transform H in (1.20), which will usually not correspond to digital filtering. We restrict attention to the case where $\{w_n\}$ form a frame for their span $\mathcal{W}$, which we call the reconstruction space. By restricting the reconstruction to the form $\hat{x} = Wd$, we are essentially imposing that the recovered signal $\hat{x}$ lie in the predefined space $\mathcal{W}$. This space can be chosen so as to lead to highly efficient interpolation methods. For example, by appropriate choice of a generator $w(t)$, the family of splines can be described using (1.19) [37, 38, 39].

1.4 Optimization methods

The fundamental problem we wish to address in this chapter is the following. Given the (noisy) samples of a signal $\tilde{c} = S^*x + u$ and some prior knowledge of the form $x \in \mathcal{A}$, produce a reconstruction $\hat{x}$ that is close to x in some sense. The set $\mathcal{A}$ incorporates our knowledge about the typical input signals and can be a subspace, as in (1.10), or an ellipsoid, as in (1.13).

Assuming that the noise u is known to be norm bounded, the samples $\tilde{c}$ together with the set $\mathcal{A}$ can be used to determine the set of feasible signals:

$$\mathcal{G} = \{x : x \in \mathcal{A}, \|S^*x - \tilde{c}\| \leq \alpha\}. \tag{1.21}$$

Thus, the unknown signal lies in $\mathcal{G}$. To find a good approximation to x in $\mathcal{G}$, it is important to notice that the reconstruction error $\|\hat{x} - x\|$ of any recovery method generally depends on the unknown signal x. This renders comparison between different methods difficult, as one method may be better than another for certain input signals and worse for others. In other words, it is generally impossible to minimize the error uniformly over $\mathcal{G}$. The same phenomenon occurs in the case where the noise u is random and the goal is to minimize the mean squared error (MSE) over the set $\mathcal{A}$ [40]. Two approaches to deal with this dependency are least-squares (LS) and worst-case (minimax) design.

In the LS strategy, the reconstruction error $\|\hat{x} - x\|$ is replaced by the error-in-samples objective $\|S^*\hat{x} - \tilde{c}\|$. This approach seeks a signal $\hat{x}$ that produces samples as close as possible to the measured samples $\tilde{c}$:

$$\hat{x}_{\mathrm{LS}} = \arg\min_{x \in \mathcal{G}} \|S^*x - \tilde{c}\|^2. \tag{1.22}$$

The objective in (1.22) is convex (quadratic) in x and therefore, if $\mathcal{G}$ is a convex set, as we assume throughout the chapter, then the problem is convex. Furthermore, the LS reconstruction admits a closed-form solution for many interesting priors. Due to its simplicity, this criterion is widely used in inverse problems in general, and in sampling in particular [11, 12]. However, it is important to note that there are situations where minimization of the error-in-samples leads to a large reconstruction error. This happens, for example, when S is such that large perturbations in x lead to small perturbations in S^*x. Therefore, this method does not guarantee a small recovery error.

An alternative to the LS approach is worst-case (or minimax) design [32, 41, 42, 14, 40]. This method attempts to control the estimation error by minimizing its largest possible value. Since x is unknown, we seek the reconstruction $\hat{x}$ that minimizes the error for the worst feasible signal:

$$\hat{x}_{\mathrm{MX}} = \arg\min_{\hat{x}} \max_{x \in \mathcal{G}} \|\hat{x} - x\|^2. \tag{1.23}$$

In contrast to (1.22), here we attempt to directly control the reconstruction error $\|x - \hat{x}\|$, which is the quantity of interest in many applications. Problem (1.23), however, is more challenging than (1.22) as we discuss next.

There are several possible approaches to solving minimax problems. In convex-concave problems we can replace the order of the minimization and the maximization [43], as incorporated in the following proposition.

Proposition 1.1. *Let $\mathcal{X}$ and $\mathcal{Y}$ be convex compact sets, and let $f(x, y)$ be a continuous function which is convex in $x \in \mathcal{X}$ for every fixed $y \in \mathcal{Y}$ and concave in $y \in \mathcal{Y}$ for every fixed $x \in \mathcal{X}$. Then,*

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} f(x, y) = \max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} f(x, y),$$

and we can replace the order of the minimization and the maximization.

There are many variants of Proposition 1.1 under weaker conditions. In particular, it is sufficient that only one of the sets will be compact, and convexity may be replaced by quasi-convexity. In the case in which the minimization is easy to solve, the problem reduces to a convex optimization problem of maximizing a concave function.

Unfortunately, since the objective in (1.23) is convex in both x and $\hat{x}$, we cannot employ Proposition 1.1 to solve it. An alternative strategy is to establish a lower bound on the objective and find a vector $\hat{x}$, which is not a function of x, that achieves it. Specifically, suppose that we show that $\max_{x\in\mathcal{G}} \|\hat{x}-x\|^2 \geq \max_{x\in\mathcal{G}} g(x)$ for all $\hat{x}$, where $g(x)$ is some function of x. Then every reconstruction $\hat{x}$, which is not a function of x, that achieves the bound is a solution. Although this approach is not constructive, when applicable, it leads to a closed-form solution to the infinite-dimensional minimax problem (1.23). The minimax problems of Sections 1.5 and 1.6 will be treated using this strategy.

Another approach to solve minimax problems is to replace the inner maximization by its dual function when strong duality holds. This will result in a minimization problem that can be combined with the outer minimization. In order to follow this method, we need to be able to establish strong duality. The inner maximization in (1.23) is a nonconvex constrained quadratic optimization problem. The nonconvexity is the result of the fact that we are maximizing a convex quadratic function, rather than minimizing it. Fortunately, we will see that under the signal priors considered here, strong duality exists in some cases [44]. In other cases, the dual problem leads to an upper bound on the objective of the inner maximization and can be used to approximate the solution. The drawback of this route is that the resulting optimization problem usually does not admit a closed-form solution and must be solved numerically. This limits this method to finite-dimensional problems. The minimax problems of Section 1.8, which involve noisy samples, will be approached using this strategy.

1.5 Subspace priors

Our first focus is on cases in which the signal $x(t)$ is known to lie in a subspace $\mathcal{A}$ spanned by frame vectors $\{a_n\}$. Given a sequence of measurements $c[n] = \langle s_n, x\rangle$, namely $c = S^*x$, where the vectors $\{s_n\}$ form a frame for the sampling space $\mathcal{S}$, our goal is to produce a reconstruction $\hat{x}$ that best approximates x in some sense.

1.5.1 Unconstrained reconstruction

We begin the discussion with the case where no constraints are imposed on the reconstruction $\hat{x}$. Interestingly, we will see that in this setting the minimax and LS solutions coincide.

1.5.1.1 Least-squares recovery

As explained in Section 1.4, in the LS strategy, the reconstruction error $\|\hat{x} - x\|$ is replaced by the error-in-samples objective $\|S^*\hat{x} - c\|$. Taking into account the prior knowledge that $x \in \mathcal{A}$, the LS recovery method can be written as

$$\hat{x}_{\mathrm{LS}} = \arg\min_{x \in \mathcal{A}} \|S^*x - c\|^2. \tag{1.24}$$

By assumption, there exists a signal $x \in \mathcal{A}$ for which $c = S^*x$. Therefore, the optimal value in (1.24) is 0. The signal x attaining this optimum is not unique if $\mathcal{A} \cap \mathcal{S}^\perp \neq \{0\}$. Indeed, suppose that x is a nonzero signal in $\mathcal{A} \cap \mathcal{S}^\perp$. Then $c[n] = \langle s_n, x\rangle = 0$ for all n and clearly x cannot be reconstructed from the measurements $c[n]$. A sufficient condition ensuring the uniqueness of the solution is that $\mathcal{A} \cap \mathcal{S}^\perp = \{0\}$ and that the Hilbert space $\mathcal{H}$ of signals can be decomposed as [45]

$$\mathcal{H} = \mathcal{A} \oplus \mathcal{S}^\perp. \tag{1.25}$$

When this condition holds, we can perfectly recover x from the samples c. This condition can be easily verified in SI spaces, as we discuss below.

Since (1.24) is defined over an infinite-dimensional Hilbert space, to solve it we do not use standard techniques such as setting the derivative of the Lagrangian to 0. Instead, we rely on the properties of the relevant spaces. Specifically, to determine the set of optimal solutions to (1.24), we express x in terms of its expansion coefficients in $\mathcal{A}$. Writing $x = \sum d[n]a_n = Ad$, the optimal sequence d is the solution to

$$\hat{d}_{\mathrm{LS}} = \arg\min_{d} \|S^*Ad - c\|^2. \tag{1.26}$$

The set of solutions to this optimization problem is given in the following theorem.

Theorem 1.1. *Every solution to (1.26) is of the form*

$$\hat{d}_{LS} = (S^*A)^\dagger c + v, \tag{1.27}$$

where v is some vector in $\mathcal{N}(S^*A)$. Furthermore, the minimal norm solution is given by $\hat{d} = (S^*A)^\dagger c$.

Before proving the theorem, we first need to verify that the pseudo-inverse is well defined. If $\mathcal{S}$ and $\mathcal{A}$ have finite dimensions, say M and N respectively, then S^*A corresponds to an $M \times N$ matrix and $(S^*A)^\dagger$ is trivially a bounded operator. However, this is not necessarily true for infinite-dimensional operators. Fortunately, the fact that S and A are synthesis operators of frames, guarantees that $(S^*A)^\dagger$ is bounded, as stated in the next proposition.

Proposition 1.2. *Let S and A be set transformations corresponding to frames $\{s_n\}$ and $\{a_n\}$ respectively. Then $(S^*A)^\dagger$ is a bounded operator.*

Proof. The proof of the proposition relies on the fact that the pseudo-inverse of an operator is well defined if its range is closed. In other words, we need to show that every Cauchy sequence $c_n \in \mathcal{R}(S^*A)$ converges to a limit in $\mathcal{R}(S^*A)$. This can be established by using the lower frame bound of S, the fact that $\mathcal{R}(A)$ and $\mathcal{R}(S^*)$ are closed, and that S^* is continuous. □

We now prove Theorem 1.1.

Proof of Theorem 1.1. To see that the set of solutions to (1.26) is given by (1.27), we substitute $\hat{d}_{\rm LS}$ of (1.27) into the objective of (1.26):

$$\begin{aligned} S^*A\hat{d}_{\rm LS} - c &= (S^*A)((S^*A)^\dagger c + v) - c \\ &= (S^*A)(S^*A)^\dagger c - c \\ &= P_{\mathcal{R}(S^*A)}c - c = 0. \end{aligned} \tag{1.28}$$

The second equality follows from the fact that $v \in \mathcal{N}(S^*A)$, the third equality follows from Lemma 1.1, and the last equality is a result of $c \in \mathcal{R}(S^*A)$ since $c = S^*x$ for some $x \in \mathcal{A}$. Therefore, every vector described by (1.27) attains an optimal value of 0 in (1.26). It is also clear from (1.28) that any vector of the form $\hat{d}_{\rm LS} + w$, where $\hat{d}_{\rm LS}$ is given by (1.27) and $w \in \mathcal{N}(S^*A)^\perp$, is not a solution to (1.26). Thus, d solves (1.26) if and only if it is of the form (1.27).

Among all solutions, the one with minimal norm is given by $\hat{d}_{\rm LS} = (S^*A)^\dagger c$. This follows from the fact that $(S^*A)^\dagger c \in \mathcal{N}(S^*A)^\perp$ by definition of the pseudo-inverse, and v lies in $\mathcal{N}(S^*A)$. The Pythagorean theorem therefore implies that $\|(S^*A)^\dagger c + v\|^2 = \|(S^*A)^\dagger c\|^2 + \|v\|^2 > \|(S^*A)^\dagger c\|^2$ for any nonzero v. □

In the sequel, we take interest only in the minimal-norm solution. From Theorem 1.1, the LS recovery method with minimal-norm amounts to applying the transformation

$$H = (S^*A)^\dagger \tag{1.29}$$

to the samples c to obtain a sequence of expansion coefficients d. This sequence is then used to synthesize $\hat{x}$ via $\hat{x} = \sum_n d[n]a_n$. Thus, $\hat{x}$ is related to x by

$$\hat{x}_{\rm LS} = Ad = AHc = A(S^*A)^\dagger c = A(S^*A)^\dagger S^*x. \tag{1.30}$$

In settings where the solution is unique, namely when $\mathcal{A} \cap \mathcal{S}^\perp = \{0\}$ and (1.25) holds, the LS strategy leads to perfect recovery of x. This has a simple geometric interpretation. It is easily verified that $\mathcal{N}(S^*A) = \mathcal{N}(A)$ in this case [45]. Consequently, $A(S^*A)^\dagger S^*$ is an oblique projection with range $\mathcal{A}$ and null space $\mathcal{S}^\perp$, denoted by $E_{\mathcal{A}\mathcal{S}^\perp}$. To see this, note that every $x \in \mathcal{A}$ can be written as $x = Ad$ for some $d \in \ell_2$ and therefore

$$A(S^*A)^\dagger S^*x = A(S^*A)^\dagger S^*Ad = AP_{\mathcal{N}(S^*A)^\perp}d = AP_{\mathcal{N}(A)^\perp}d = Ad = x. \tag{1.31}$$

On the other hand, for any $x \in \mathcal{S}^\perp$ we have $A(S^*A)^\dagger S^*x = 0$. Thus, $\hat{x} = E_{\mathcal{A}\mathcal{S}^\perp}x = x$ for any $x \in \mathcal{A}$, which implies that the LS reconstruction (1.30) coinsides with

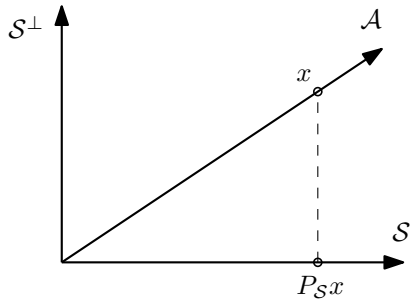

Figure 1.5: A unique element in $\mathcal{A}$ which is consistent with the samples in $\mathcal{S}$ can be recovered from the known samples.

the original signal x. As a special case, if $\mathcal{A} = \mathcal{S}$ then the LS solution reduces to an orthogonal projection $P_{\mathcal{S}} = S(S^*S)^\dagger S^*$.

The geometric explanation of $\hat{x} = E_{\mathcal{A}\mathcal{S}^\perp}x$ follows from the fact that knowing the samples c is equivalent to knowing the orthogonal projection of the signal onto the sampling space, since $P_{\mathcal{S}}x = S(S^*S)^\dagger S^*x = S(S^*S)^\dagger c$. The direct-sum condition (1.25) ensures that there is a unique vector in $\mathcal{A}$ with the given projection onto $\mathcal{S}$. As depicted in Fig. 1.5, 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. Therefore, under (1.25), x can be perfectly recovered from c by using H of (1.29). To conclude, we see that if $\mathcal{A} \cap \mathcal{S}^\perp = \{0\}$ then the LS approach leads to perfect recovery. If the intersection is non-trivial, on the other hand, then clearly perfect reconstruction is not possible since there are generally infinitely many signals in $\mathcal{A}$ yielding the same samples $c = S^*x$.

We now turn our attention to the case in which $\mathcal{A}$ and $\mathcal{S}$ are SI spaces with generators $a(t)$ and $s(t)$ respectively. In this setting, the operator S^*A corresponds to convolution with the sequence $r_{SA}[n] = (a(t) * s(-t))(n)$. To see this, let $c = S^*Ad$. From Definition 1.1,

$$\begin{aligned} c[k] &= \int_{-\infty}^{\infty} s(t-k) \sum_{n=-\infty}^{\infty} d[n]a(t-n)dt \\ &= \sum_{n=-\infty}^{\infty} d[n](a(t) * s(-t))(k-n) \\ &= (d[n] * r_{SA}[n])[k]. \end{aligned} \tag{1.32}$$

Therefore, the correction $H = (S^*A)^\dagger$ is a digital filter with frequency response

$$H(e^{j\omega}) = \begin{cases} \frac{1}{\phi_{SA}(e^{j\omega})}, & \phi_{SA}\left(e^{j\omega}\right) \neq 0; \\ 0, & \phi_{SA}\left(e^{j\omega}\right) = 0, \end{cases} \tag{1.33}$$

where $\phi_{SA}(e^{j\omega})$ is the DTFT of $r_{SA}[n]$, and is given by

$$\phi_{SA}(e^{j\omega}) = \sum_{k=-\infty}^{\infty} S^*(\omega + 2\pi k)A(\omega + 2\pi k). \tag{1.34}$$

The overall scheme fits that depicted in Fig. 1.4, with $w(t) = a(t)$ and $h[n]$ given by (1.33).

The direct sum condition (1.25) ensuring perfect recovery, can be verified easily in SI spaces. Specifically, (1.25) is satisfied if and only if [46] the supports $\mathcal{I}_\mathcal{A}$ and $\mathcal{I}_\mathcal{S}$ of $\phi_{SS}(e^{j\omega})$ and $\phi_{AA}(e^{j\omega})$ respectively coincide, and there exists a constant $\alpha > 0$ such that $|\phi_{SA}(e^{j\omega})| > \alpha$, for all $\omega \in \mathcal{I}_\mathcal{A}$.

We conclude the discussion with a non-intuitive example in which a signal that is not bandlimited is filtered with a LPF prior to sampling, and still can be perfectly reconstructed from the resulting samples.

Example 1.1: Consider a signal $x(t)$ formed by exciting an RC circuit with a modulated impulse train $\sum_n d[n]\delta(t-n)$, as shown in Fig. 1.6(a). 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

$$x(t) = \frac{1}{\tau}\sum_{n=-\infty}^{\infty} d[n]\exp\{-(t-n)/\tau\}u(t-n). \tag{1.35}$$

Clearly, $x(t)$ is not bandlimited. Now, suppose that $x(t)$ is filtered by an ideal LPF $s(t) = \mathrm{sinc}(t)$ and then sampled at times $t = n$ to obtain the sequence $c[n]$. The signal $x(t)$ and its samples are depicted in Fig. 1.6(b). 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, as shown in Fig. 1.6(c). However, it is easily verified that if $\tau < \pi^{-1}$, then $|\phi_{SA}(e^{j\omega})| > (1-\pi^2\tau^2)^{-1/2} > 0$ for all $\omega \in [-\pi, \pi]$ so that condition (1.25) is satisfied. Therefore, perfect recovery is possible in this setup using the LS approach. The digital correction filter (1.33) in this case can be shown to be

$$h[n] = \begin{cases} 1 & n = 0; \\ \frac{\tau}{n}(-1)^n & n \neq 0. \end{cases} \tag{1.36}$$

Thus, to reconstruct $x(t)$ we need to excite an identical RC circuit with an impulse train modulated by the sequence $d[n] = h[n] * c[n]$. The entire sampling-reconstruction setup is depicted in Fig. 1.6(a).

1.5.1.2 Minimax recovery

We now treat the recovery of x via a minimax framework:

$$\hat{x}_{\mathrm{MX}} = \arg\min_{\hat{x}} \max_{x\in\mathcal{G}} \|\hat{x} - x\|^2, \tag{1.37}$$

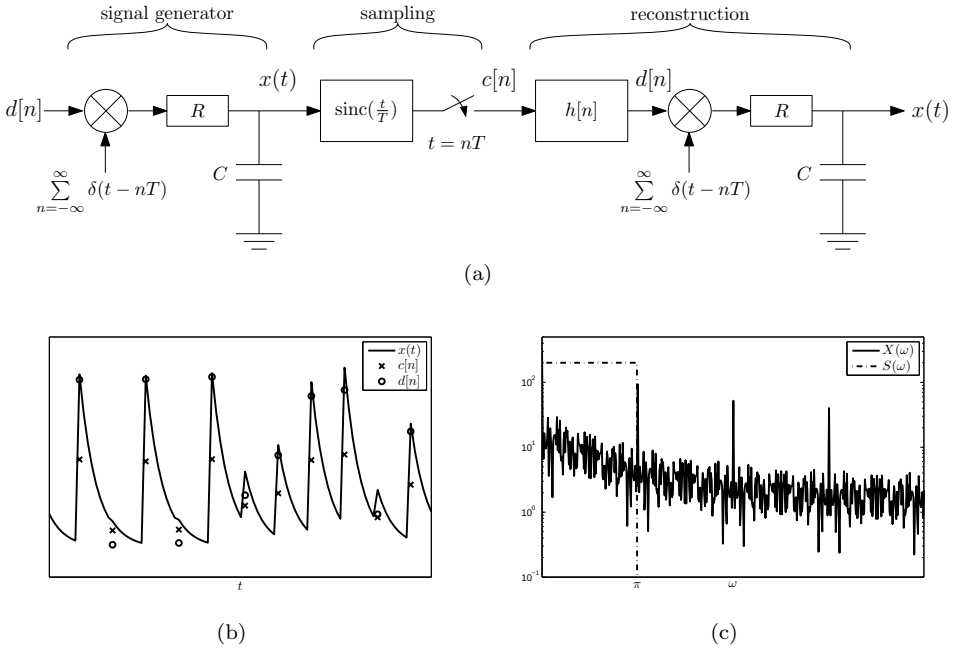

Figure 1.6: A non-bandlimited signal $x(t)$, formed by exciting an RC-circuit with a modulated impulse train, is sampled after passing through an ideal LPF and then perfectly reconstructed using the LS approach. (a) Sampling and reconstruction setup. (b) The signal $x(t)$, its samples $c[n]$ and expansion coefficients $d[n]$. (c) The signal $X(\omega)$ and the sampling filter $S(\omega)$.

where $\mathcal{G}$ is the set of signals $x \in \mathcal{A}$ satisfying $S^*x = c$.

To approach this problem, notice that $\hat{x}_{\mathrm{MX}}$ must lie in $\mathcal{A}$, as any $\hat{x} \notin \mathcal{A}$ can be improved upon by projecting it onto $\mathcal{A}$: $\|\hat{x} - x\|^2 \geq \|P_{\mathcal{A}}\hat{x} - x\|^2$ for any $\hat{x}$ and $x \in \mathcal{A}$. Therefore, we can express both $\hat{x}$ and x in terms of their expansion coefficients in $\mathcal{A}$, by writing $\hat{x} = A\hat{d}$ and $x = Ad$. To guarantee that the error in (1.37) cannot grow without bound, the sequence d should be constrained to lie in some bounded set. We therefore impose the additional requirement that $\|d\| \leq \rho$, for some $\rho > 0$. Problem (1.37) can then be reformulated as

$$\min_{\hat{d}} \max_{d \in \mathcal{D}} \|A\hat{d} - Ad\|^2, \tag{1.38}$$

where $\mathcal{D} = \{d : S^*Ad = c, \|d\| \leq \rho\}$. As we now show, the choice of ρ does not affect the solution, as long as $\mathcal{D}$ is a nonempty set.

Theorem 1.2. *A solution to problem (1.38) is* $\hat{d} = (S^*A)^\dagger c$.

Proof. As we have seen in the proof of Theorem 1.1, assuming that $c \in \mathcal{R}(S^*A)$, a sequence d satisfies $S^*Ad = c$ if and only if it is of the form $d = (S^*A)^\dagger c + v$,

where v is some vector in $\mathcal{N}(S^*A)$. Furthermore, $(S^*A)^\dagger c \in \mathcal{N}(S^*A)^\perp$ so that $\|v\|^2 = \|d\|^2 - \|(S^*A)^\dagger c\|^2$. Therefore, the inner maximization in (1.38) becomes

$$\|A(\hat{d} - (S^*A)^\dagger c)\|^2 + \max_{v\in\mathcal{V}}\{\|Av\|^2 - 2v^*A^*A(\hat{d} - (S^*A)^\dagger c)\}, \tag{1.39}$$

where

$$\mathcal{V} = \left\{v : v \in \mathcal{N}(S^*A), \|v\|^2 \le \rho^2 - \|(S^*A)^\dagger c\|^2\right\}. \tag{1.40}$$

Since $\mathcal{V}$ is a symmetric set, the vector v attaining the maximum in (1.39) must satisfy $v^*A^*A(\hat{d} - (S^*A)^\dagger c) \le 0$, as we can change the sign of v without effecting the constraint. Therefore,

$$\max_{v\in\mathcal{V}}\{\|Av\|^2 - 2v^*A^*A(\hat{d} - (S^*A)^\dagger c)\} \ge \max_{v\in\mathcal{V}}\|Av\|^2. \tag{1.41}$$

Combining (1.41) and (1.39) we have that

$$\begin{aligned}\min_{\hat{d}}\max_{d\in\mathcal{D}}\|A\hat{d} - Ad\|^2 &\ge \min_{\hat{d}}\{\|A(\hat{d} - (S^*A)^\dagger c)\|^2 + \max_{v\in\mathcal{V}}\|Av\|^2\}\\ &= \max_{v\in\mathcal{V}}\|Av\|^2,\end{aligned} \tag{1.42}$$

where the equality is a result of solving the minimization, which is obtained *e.g.,* at $\hat{d} = (S^*A)^\dagger c$.

We now show that the inequality in (1.42) can be achieved with $\hat{d} = (S^*A)^\dagger c$. Indeed, substituting this choice of $\hat{d}$ in (1.39), we have that

$$\max_{d\in\mathcal{D}}\|A\hat{d} - Ad\|^2 = \max_{v\in\mathcal{V}}\{\|Av\|^2 - 2v^*A(\hat{d} - (S^*A)^\dagger c)\} = \max_{v\in\mathcal{V}}\|Av\|^2, \tag{1.43}$$

concluding the proof. □

We conclude that a solution to the minimax problem (1.37) is given by

$$\hat{x}_{\mathrm{MX}} = A(S^*A)^\dagger c, \tag{1.44}$$

coinciding with the LS solution (1.30). We also see that, as in the LS strategy, the expansion coefficients of the recovery $\hat{x}$ in $\mathcal{A}$ are obtained by applying $H = (S^*A)^\dagger$ on the samples c.

Although the minimax and LS approaches coincide in the unconstrained subspace setting discussed thus far, we will wee that these strategies lead to quite different reconstruction methods when the reconstruction process is constrained. In Section 1.6 we will also show that the results differ under a smoothness prior.

1.5.2 Constrained reconstruction

Up until now we specified the sampling process but did not restrict the reconstruction or interpolation kernel $w(t)$ in Fig. 1.4. We now address the problem of approximating x using a predefined set of reconstruction functions $\{w_n\}$, which form a frame for the reconstruction space $\mathcal{W}$. Given sampling functions $\{s_n\}$ and a fixed set of reconstruction functions $\{w_n\}$ an important question is how

to design the correction transform H so that the output $\hat{x}$ is a good approximation of the input signal x in some sense. To handle this problem, we extend the two approaches discussed in Section 1.5.1 to the constrained setup. However, in contrast to the previous section, where perfect recovery was under a direct sum assumption, here $\hat{x}$ must lie in the space $\mathcal{W}$. Therefore, if x does not lie in $\mathcal{W}$ to begin with, then $\hat{x}$ cannot be equal x.

1.5.2.1 Least squares recovery

To obtain a reconstruction in $\mathcal{W}$ within the LS methodology we reformulate (1.24) as

$$\hat{x}_{\mathrm{CLS}} = \arg\min_{x\in\mathcal{W}} \|S^*x - c\|^2. \tag{1.45}$$

Namely, the reconstruction $\hat{x}_{\mathrm{CLS}} \in \mathcal{W}$ should yield samples as close as possible to the measured sequence c. Note that (1.45) actually ignores our prior knowledge that $x \in \mathcal{A}$, a problem which is inevitable when working with the error-in-samples criterion.

Problem (1.45) is similar to (1.24) with $\mathcal{A}$ replaced by $\mathcal{W}$. However, here c does not necessarily lie in $\mathcal{R}(S^*W)$. Therefore, there does not necessarily exist an $x \in \mathcal{W}$ giving rise to the measured samples c, and consequently the minimal distance is generally not 0.

Theorem 1.3. *A solution to (1.45) is* $\hat{x}_{CLS} = W(S^*W)^\dagger c$.

Proof. Let $\hat{d}$ denote the expansion coefficients of the reconstruction, so that $\hat{x} = W\hat{d}$, and let $\hat{c} = S^*W\hat{d}$ be the samples it produces. Then, (1.45) can be written as

$$\min_{\hat{c}\in\mathcal{R}(S^*W)} \|\hat{c} - c\|^2. \tag{1.46}$$

This formulation shows that the optimal $\hat{c}$ is the projection of c onto $\mathcal{R}(S^*W)$:

$$\hat{c} = S^*W\hat{d} = P_{\mathcal{R}(S^*W)}c = (S^*W)(S^*W)^\dagger c, \tag{1.47}$$

from which the result follows. □

The solution of Theorem 1.3 has the same structure as the unconstrained LS reconstruction (1.30) with A replaced by W. Furthermore, as in Section 1.5.1.1, this solution is not unique if $\mathcal{W} \cap \mathcal{S}^\perp \neq \{0\}$.

It is interesting to study the relation between the unconstrained and constrained solutions. As we have seen, $\hat{x}_{\mathrm{LS}}$ of (1.30) is consistent, namely $S^*\hat{x}_{\mathrm{LS}} = c$. Therefore, $\hat{x}_{\mathrm{CLS}}$ can be expressed in terms of $\hat{x}_{\mathrm{LS}}$:

$$\hat{x}_{\mathrm{CLS}} = W(S^*W)^\dagger c = W(S^*W)^\dagger S^*\hat{x}_{\mathrm{LS}}. \tag{1.48}$$

The geometric meaning of this relation is best understood when $\mathcal{H} = \mathcal{W} \oplus \mathcal{S}^\perp$. Then, $\hat{x}_{\rm CLS}$ is the oblique projection of $\hat{x}_{\rm LS}$ onto $\mathcal{W}$ along $\mathcal{S}^\perp$:

$$\hat{x}_{\rm CLS} = E_{\mathcal{W}\mathcal{S}^\perp}\hat{x}_{\rm LS}. \tag{1.49}$$

Figure 1.7 depicts $\hat{x}_{\rm LS}$ and $\hat{x}_{\rm CLS}$ in a situation where $\mathcal{A}$ and $\mathcal{S}^\perp$ satisfy the direct-sum condition (1.25) so that $\hat{x}_{\rm LS} = \hat{x}_{\rm MX} = x$, and also $\mathcal{H} = \mathcal{W} \oplus \mathcal{S}^\perp$ implying that (1.49) holds. This example highlights the disadvantage of the LS formulation. In this setting we are constrained to yield $\hat{x} \in \mathcal{W}$. But since x can be determined from the samples c in this case, so can its best approximation in $\mathcal{W}$, which is given by $P_{\mathcal{W}}x = P_{\mathcal{W}}\hat{x}_{\rm LS}$. This alternative is also shown in Fig. 1.7, and is clearly advantageous to $\hat{x}_{\rm CLS}$ in terms of squared error *for every* x. We will see in Section 1.5.2.2 that orthogonally projecting $\hat{x}_{\rm LS}$ onto the reconstruction space $\mathcal{W}$ can be motivated also when condition (1.25) does not hold.

The constrained LS method can be easily implemented in situations where $\mathcal{W}$ and $\mathcal{S}$ are SI spaces with generators $w(t)$ and $s(t)$ respectively. As we have seen, the operator S^*W corresponds to convolution with the sequence $r_{SW}[n] = (w(t) * s(-t))(n)$ and thus the correction transform $H = (S^*W)^\dagger$ is a digital filter whose frequency response is

$$H(e^{j\omega}) = \begin{cases} \frac{1}{\phi_{SW}(e^{j\omega})}, & \phi_{SW}\left(e^{j\omega}\right) \neq 0; \\ 0, & \phi_{SW}\left(e^{j\omega}\right) = 0, \end{cases} \tag{1.50}$$

where $\phi_{SW}(e^{j\omega})$ is the DTFT of $r_{SW}[n]$, which is given by (1.34) with $A(\omega)$ replaced by $W(\omega)$. To conclude, reconstruction is performed by the scheme depicted in Fig. 1.4, where the reconstruction kernel is $w(t)$ and the digital correction filter is given by (1.50).

1.5.2.2 Minimax recovery

We now treat the constrained recovery setting via a worst-case design strategy. The constraint $\hat{x} \in \mathcal{W}$ leads to an inherent limitation on the minimal achievable reconstruction error. Indeed, since $x \in \mathcal{A}$, the reconstruction error cannot be 0 unless $\mathcal{W} \subseteq \mathcal{A}$ [32]. From (1.4), we know that the best approximation in $\mathcal{W}$ of any signal x is given by $\hat{x} = P_{\mathcal{W}}x$, which in general cannot be computed from the sequence of samples $c[n]$. Therefore, we consider here the minimization of the *regret*, which is defined by $\|\hat{x} - P_{\mathcal{W}}x\|^2$. Since the regret is a function of the unknown signal x, we seek the reconstruction $\hat{x} \in \mathcal{W}$ minimizing the worst-case regret [47, 48, 32, 14]. Our problem is thus

$$\hat{x}_{\rm CMX} = \min_{\hat{x}\in\mathcal{W}} \max_{x\in\mathcal{G}} \|\hat{x} - P_{\mathcal{W}}x\|^2, \tag{1.51}$$

where $\mathcal{G}$ is the set of signals $x \in \mathcal{A}$ satisfying $S^*x = c$.

To solve (1.51) we express $\hat{x}$ and x in terms of their expansion coefficients in $\mathcal{W}$ and $\mathcal{A}$ respectively, by writing $\hat{x} = W\hat{d}$ and $x = Ad$. As in the unconstrained setting, we require that $\|d\| \leq \rho$ for some $\rho > 0$, in order for the inner maximization

to be bounded. Therefore, problem (1.51) can be written as

$$\min_{\hat{d}} \max_{d\in\mathcal{D}} \|W\hat{d} - P_{\mathcal{W}}Ad\|^2, \tag{1.52}$$

where $\mathcal{D} = \{d : S^*Ad = c, \|d\| \le \rho\}$.

Theorem 1.4. *A solution to problem (1.51) is* $\hat{d} = (W^*W)^\dagger W^*A(S^*A)^\dagger c$.

Proof. The set $\mathcal{D}$ consists of all sequences of the form $d = (S^*A)^\dagger c + v$, where v is some vector in $\mathcal{N}(S^*A)$, with $\|v\|^2 \le \|d\|^2 - \|(S^*A)^\dagger c\|^2$. Therefore, the inner maximization in (1.52) becomes

$$\|W\hat{d} - P_{\mathcal{W}}A(S^*A)^\dagger c\|^2 + \max_{v\in\mathcal{V}}\{\|P_{\mathcal{W}}Av\|^2 - 2(P_{\mathcal{W}}Av)^*(W\hat{d} - P_{\mathcal{W}}A(S^*A)^\dagger c)\}, \tag{1.53}$$

where $\mathcal{V}$ is given by (1.40). Since $\mathcal{V}$ is a symmetric set, the vector v attaining the maximum in (1.53) must satisfy $(P_{\mathcal{W}}Av)^*(W\hat{d} - P_{\mathcal{W}}(S^*A)^\dagger c) \le 0$, as we can change the sign of v without effecting the constraint. Consequently,

$$\max_{v\in\mathcal{V}}\{\|P_{\mathcal{W}}Av\|^2 - 2(P_{\mathcal{W}}Av)^*(W\hat{d} - P_{\mathcal{W}}A(S^*A)^\dagger c)\} \ge \max_{v\in\mathcal{V}} \|P_{\mathcal{W}}Av\|^2. \tag{1.54}$$

Combining (1.54) and (1.53) we have that

$$\begin{aligned}\min_{\hat{d}} \max_{d\in\mathcal{D}} \|W\hat{d} - P_{\mathcal{W}}Ad\|^2 &\ge \min_{\hat{d}}\{\|W\hat{d} - P_{\mathcal{W}}A(S^*A)^\dagger c\|^2 + \max_{v\in\mathcal{V}} \|P_{\mathcal{W}}Av\|^2\} \\ &= \max_{v\in\mathcal{V}} \|P_{\mathcal{W}}Av\|^2,\end{aligned} \tag{1.55}$$

where the equality is a result of solving the minimization, which is obtained *e.g.,* at

$$\hat{d} = (W^*W)^\dagger W^*A(S^*A)^\dagger c. \tag{1.56}$$

We now show that the inequality can be achieved with $\hat{d}$ given by (1.56). Substituting this into (1.53), we have that

$$\begin{aligned}\max_{d\in\mathcal{D}} \|W\hat{d} - P_{\mathcal{W}}Ad\|^2 &= \max_{v\in\mathcal{V}}\{\|P_{\mathcal{W}}Av\|^2 - 2(P_{\mathcal{W}}Av)^*(W\hat{d} - P_{\mathcal{W}}A(S^*A)^\dagger c)\} \\ &= \max_{v\in\mathcal{V}} \|P_{\mathcal{W}}Av\|^2,\end{aligned} \tag{1.57}$$

from which the proof follows. □

We conclude that the solution to the minimax problem (1.51) is given by

$$\hat{x}_{\mathrm{CMX}} = W\hat{d} = W(W^*W)^\dagger W^*A(S^*A)^\dagger c = P_{\mathcal{W}}A(S^*A)^\dagger c. \tag{1.58}$$

In contrast to the constrained LS reconstruction of Theorem 1.3, the minimax regret solution of Theorem 1.4 explicitly depends on A. Hence, the prior knowledge that $x \in \mathcal{A}$ plays a role, as one would expect. It is also readily observed that the relation between the unconstrained and constrained minimax solutions is different than in the LS approach. Identifying in (1.58) the expression

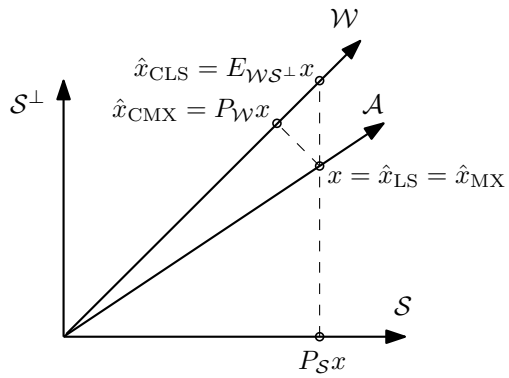

Figure 1.7: When condition (1.25) holds, $x = x_{\rm LS} = x_{\rm MX}$ and thus the signal $x \in \mathcal{A}$ can be recovered from the samples $c[n]$, allowing to compute its projection onto $\mathcal{W}$. The constrained minimax approach indeed yields $\hat{x}_{\rm CMX} = P_{\mathcal{W}} x$ whereas the constrained least-squares criterion leads to $\hat{x}_{\rm CLS} = E_{\mathcal{W}\mathcal{S}^\perp} x$.

$A(S^*A)^\dagger c = \hat{x}_{\rm MX} = \hat{x}_{\rm LS}$, the constrained minimax recovery can be written as

$$\hat{x}_{\rm CMX} = P_{\mathcal{W}} \hat{x}_{\rm MX}, \tag{1.59}$$

so that the constrained solution is the orthogonal projection onto $\mathcal{W}$ of the unconstrained reconstruction. In Section 1.5.2.1 we discussed the superiority of this approach in situations where the spaces $\mathcal{S}$ and $\mathcal{A}$ satisfy the direct-sum condition (1.25), as shown in Fig. 1.7. We now see that this strategy stems from the minimization of the worst case regret, for *any two spaces* $\mathcal{S}$ and $\mathcal{A}$.

Let us now examine the case where $\mathcal{S}$, $\mathcal{A}$, and $\mathcal{W}$ are SI spaces with generators $s(t)$, $a(t)$, and $w(t)$ respectively. As shown in Section 1.5.1.1, each of the operators $(W^*W)^\dagger$, W^*A and $(S^*A)^\dagger$ corresponds to a digital filter. Therefore, the overall reconstruction scheme is that depicted in Fig. 1.4 with a digital correction filter $H(e^{j\omega})$ given by

$$H\left(e^{j\omega}\right) = \begin{cases} \frac{\phi_{WA}(e^{j\omega})}{\phi_{SA}(e^{j\omega})\phi_{WW}(e^{j\omega})}, & \phi_{SA}(e^{j\omega})\phi_{WW}(e^{j\omega}) \neq 0; \\ 0, & \phi_{SA}(e^{j\omega})\phi_{WW}(e^{j\omega}) = 0, \end{cases} \tag{1.60}$$

where $\phi_{WA}(e^{j\omega})$, $\phi_{SA}(e^{j\omega})$ and $\phi_{WW}(e^{j\omega})$ follow from (1.34) with the corresponding substitution of the filters $W(\omega)$, $A(\omega)$ and $S(\omega)$.

To demonstrate the minimax regret recovery procedure, we now revisit Example 1.1 imposing a constraint on the recovery mechanism.

Example 1.2: Suppose that the signal $x(t)$ of (1.35) is sampled at the integers after passing through the anti-aliasing filter $s(t) = \mathrm{sinc}(t)$, as in Example 1.1. We would now like to recover $x(t)$ from the samples $c[n]$ using a standard zero-order-hold digital-to-analog convertor. The corresponding reconstruction filter is

therefore $w(t) = u(t) - u(t-1)$, where $u(t)$ is the unit step function. To compute the digital compensation filter (1.60), we note that $\phi_{WW}(e^{j\omega}) = 1$ in our case. Furthermore, the filter $1/\phi_{SA}(e^{j\omega})$ is given by (1.36), as we have already seen in Example 1.1. It can be easily shown that the remaining term, $\phi_{WA}(e^{j\omega})$, corresponds to the filter

$$h_{WA}[n] = \begin{cases} e^{\frac{n}{\tau}}\left(1 - e^{-\frac{1}{\tau}}\right), & n \leq 0; \\ 0, & n > 0. \end{cases} \tag{1.61}$$

Therefore, the sequence $d[n]$ feeding the DAC is obtained by convolving the samples $c[n]$ with $h[n]$ of (1.36) and then by $h_{WA}[n]$ of (1.61).

To summarize, we have seen that treating the constrained reconstruction scenario within the minimax regret framework leads to a simple and plausible recovery method. In contrast, the constrained LS approach does not take the prior into account and is thus often inferior in terms of squared error in this setting.

1.6 Smoothness priors

We now treat the problem of approximating $x(t)$ from its samples, based on the knowledge that it is smooth. Specifically, here x is assumed to obey (1.13) with some $\rho > 0$.

1.6.1 Unconstrained reconstruction

1.6.1.1 Least-squares approximation

We begin by approximating a smooth signal x via the minimization of the error-in-samples criterion. To take the smoothness prior into account, we define the set $\mathcal{G}$ of feasible signals as $\mathcal{G} = \{x : \|Lx\| \leq \rho\}$. The LS problem is then

$$\hat{x}_{\mathrm{LS}} = \arg\min_{x \in \mathcal{G}} \|S^* x - c\|^2. \tag{1.62}$$

Since, by assumption, there exists an x in $\mathcal{G}$ giving rise to the measured samples c, the optimal value in (1.62) is 0. Furthermore, there may be infinitely many solutions in $\mathcal{G}$ yielding 0 error-in-samples, as demonstrated in Fig. 1.8(b). In this figure, the solid vertical segment is the set of signals satisfying $S^* x = c$ and $\|Lx\| \leq \rho$. To resolve this ambiguity, we seek the smoothest reconstruction among all possible solutions:

$$\hat{x}_{\mathrm{LS}} = \arg\min_{x \in \mathcal{G}} \|Lx\|, \tag{1.63}$$

where now $\mathcal{G} = \{x : S^* x = c\}$.

Problem (1.63) is a linearly constrained quadratic program with a convex objective. In finite dimensions there always exists a solution to this kind of

problem. However, in infinite dimensions this is no longer true [49, Chapter 11]. To guarantee the existence of a solution we focus on situations in which the operator L^*L is bounded from above and below, so that there exist constants $0 < \alpha_L \leq \beta_L < \infty$ such that

$$\alpha_L \|x\|^2 \leq \|L^*Lx\|^2 \leq \beta_L \|x\|^2. \tag{1.64}$$

Since L^*L is Hermitian, this condition also implies that $(L^*L)^{-1}$ is bounded and that $\beta_L^{-1}\|x\|^2 \leq \|(L^*L)^{-1}x\|^2 \leq \alpha_L^{-1}\|x\|^2$ for any $x \in \mathcal{H}$.

Theorem 1.5. *Assume that the operator L satisfies condition (1.64). Then the solution (1.63) is given by*

$$\hat{x}_{LS} = \tilde{W}(S^*\tilde{W})^\dagger c, \tag{1.65}$$

where

$$\tilde{W} = (L^*L)^{-1}S. \tag{1.66}$$

Proof. Since $(L^*L)^{-1}$ is upper- and lower-bounded, and S satisfies the frame condition (1.7), $\tilde{W}$ is a synthesis operator of a frame. From Proposition 1.2, it then follows that $(S^*\tilde{W})^\dagger$ is bounded.

To solve (1.63), define the operator

$$E = \tilde{W}(S^*\tilde{W})^\dagger S^*, \tag{1.67}$$

so that $\hat{x}_{\mathrm{LS}}$ of (1.65) is given by $\hat{x}_{\mathrm{LS}} = Ex$. Now, any x can be decomposed as

$$x = Ex + (I - E)x = Ex + v, \tag{1.68}$$

where $v = (I - E)x$. In addition,

$$S^*E = S^*(L^*L)^{-1}S(S^*(L^*L)^{-1}S)^\dagger S^* = P_{\mathcal{N}(S^*(L^*L)^{-1}S)^\perp}S^* = S^*, \tag{1.69}$$

where we used the fact that $\mathcal{N}(S^*(L^*L)^{-1}S) = \mathcal{N}(S) = \mathcal{R}(S^*)^\perp$. Therefore, $S^*x = c$ and $S^*v = 0$. Next, using the identity $U^\dagger U U^\dagger = U^\dagger$, it can be verified that

$$E^*L^*L(I - E) = 0. \tag{1.70}$$

Consequently,

$$\|Lx\|^2 = \|LEx\|^2 + \|L(I - E)x\|^2 = \|LEx\|^2 + \|Lv\|^2, \tag{1.71}$$

and thus $\|Lx\|^2$ is minimized by choosing $Lv = 0$. Finally, if $S^*x = c$ then $Ex = \hat{x}_{\mathrm{LS}}$ of (1.65). □

In Section 1.5, we have seen that (1.65) corresponds to the LS and minimax reconstructions when we have prior knowledge that x lies in the range of $\tilde{W}$, which we denote by $\tilde{\mathcal{W}}$. Thus, this approach can be viewed as first determining the optimal reconstruction space given by (1.66), and then computing the LS (or minimax) reconstruction under the subspace prior $x \in \tilde{\mathcal{W}}$.

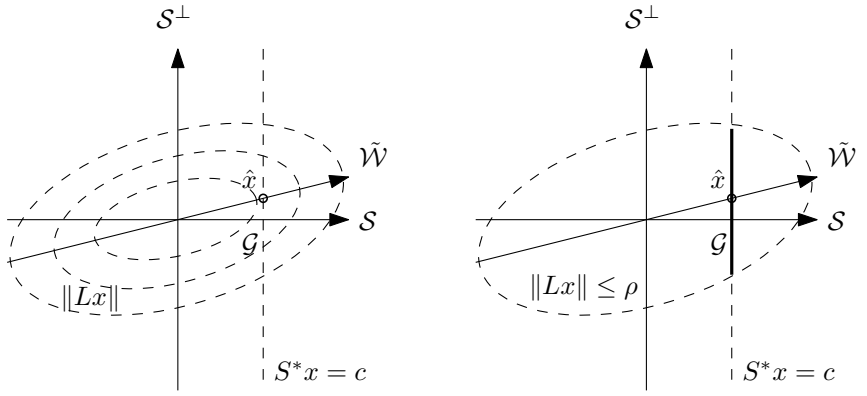

(a) LS (b) Minimax

Figure 1.8: Geometric interpretation of the LS (a) and minimax (b) recoveries.

Figure 1.8(a) shows a geometric interpretation of the LS solution. The set of feasible signals is the subspace $S^*x = c$, which is orthogonal to $\mathcal{S}$ (vertical dashed line). The dashed ellipsoids are the level sets of the objective function $\|Lx\|$. The LS solution is the intersection between the vertical line and the ellipsoid for which it constitutes a tangent. The reconstruction space $\tilde{\mathcal{W}}$, is the line connecting all possible reconstructions (for all possible sample sequences c).

As a special case, we may choose to produce the minimal-norm consistent reconstruction $\hat{x}$ by letting L be the identity operator I. This leads to $\tilde{\mathcal{W}} = \mathcal{S}$ and consequently, $\hat{x}$ is the orthogonal projection onto the sampling space, $\hat{x} = S(S^*S)^\dagger S^*x = P_{\mathcal{S}}x$. This can also be seen by noting that any reconstruction $\hat{x}$ which yields the samples c has the form $\hat{x} = P_{\mathcal{S}}x + v$ where v is an arbitrary vector in $\mathcal{S}^\perp$. The minimal-norm approximation corresponds to the choice $v = 0$.

If L is an LTI operator corresponding to convolution with a kernel whose CTFT is $L(\omega)$, then $(L^*L)^{-1}$ corresponds to filtering with $1/|L(\omega)|^2$. In this case, if the sampling space $\mathcal{S}$ is SI, then $\tilde{\mathcal{W}}$ is a SI space with generator $\tilde{w}(t)$ whose CTFT is $\tilde{W}(\omega) = S(\omega)/|L(\omega)|^2$. This is because the nth reconstruction function $w_n(t)$ is a filtered version of the corresponding sampling function $s_n(t) = s(t-n)$, namely $\tilde{W}_n(\omega) = S_n(\omega)/|L(\omega)|^2$. As shown in the previous section, the correction transform H yielding the expansion coefficients $d[n]$ is also LTI in this case, *i.e.,* it corresponds to digital filtering. Therefore, the overall reconstruction scheme is that shown in Fig. 1.4 where now the reconstruction kernel is

$$\tilde{W}(\omega) = \frac{S(\omega)}{|L(\omega)|^2}, \tag{1.72}$$

and the digital correction filter is

$$H(e^{j\omega}) = \begin{cases} \frac{1}{\phi_{S\tilde{W}}(e^{j\omega})}, & \phi_{S\tilde{W}}\left(e^{j\omega}\right) \neq 0; \\ 0, & \phi_{S\tilde{W}}\left(e^{j\omega}\right) = 0, \end{cases} \tag{1.73}$$

Here, the filter $\phi_{S\tilde{W}}(e^{j\omega})$ follows from (1.34) with $A(\omega)$ replaced by $\tilde{W}(\omega)$.

1.6.1.2 Minimax recovery

We now treat the problem of reconstructing a smooth signal from its samples via a worst-case design approach. The prior information we have can be used to construct a set $\mathcal{G}$ of all possible input signals:

$$\mathcal{G} = \{x : S^*x = c, \|Lx\| \leq \rho\}. \tag{1.74}$$

As in Section 1.6.1.1, $\rho > 0$ is a assumed to be large enough so that $\mathcal{G}$ is nonempty. 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 $\mathcal{G}$:

$$\hat{x}_{\mathrm{MX}} = \min_{\hat{x}} \max_{x\in\mathcal{G}} \|\hat{x} - x\|^2. \tag{1.75}$$

Theorem 1.6. *The solution to problem (1.75) coincides with the LS approach, namely $\hat{x}_{MX}$ equals $\hat{x}_{LS}$ of (1.65).*

Proof. For any signal x satisfying the consistency constraint $S^*x = c$, the norm $\|Lx\|^2$ is given by (1.71), with E of (1.67). Therefore, we can write the inner maximization in (1.75) as

$$\|\hat{x} - \tilde{W}(S^*\tilde{W})^\dagger c\|^2 + \max_{v\in\mathcal{V}} \left\{\|v\|^2 - 2(\hat{x} - \tilde{W}(S^*\tilde{W})^\dagger c)^* v\right\}, \tag{1.76}$$

where

$$\mathcal{V} = \left\{v : \|Lv\|^2 \leq \rho^2 - \|L\tilde{W}(S^*\tilde{W})^\dagger c\|^2\right\}. \tag{1.77}$$

Clearly, at the maximum value of v we have that $(\hat{x} - \tilde{W}(S^*\tilde{W})^\dagger c)^* v \leq 0$ since we can change the sign of v without effecting the constraint. Therefore,

$$\max_{v\in\mathcal{V}} \left\{\|v\|^2 - 2(\hat{x} - \tilde{W}(S^*\tilde{W})^\dagger c)^* v\right\} \geq \max_{v\in\mathcal{V}} \|v\|^2. \tag{1.78}$$

Combining (1.78) and (1.76),

$$\min_{\hat{x}} \max_{x\in\mathcal{G}} \|\hat{x} - x\|^2 \geq \min_{\hat{x}} \left\{\|\hat{x} - \tilde{W}(S^*\tilde{W})^\dagger c\|^2 + \max_{v\in\mathcal{V}} \|v\|^2\right\} = \max_{v\in\mathcal{V}} \|v\|^2, \tag{1.79}$$

where the equality is a result of solving the inner minimization, obtained at $\hat{x} = \tilde{W}(S^*\tilde{W})^\dagger c$. We now show that the inequality can be achieved with $\hat{x} = \tilde{W}(S^*\tilde{W})^\dagger c$. Indeed, with this choice of $\hat{x}$, (1.76) implies that

$$\max_{x\in\mathcal{G}} \|\hat{x} - x\|^2 = \max_{v\in\mathcal{V}} \left\{\|v\|^2 - 2(\hat{x} - \tilde{W}(S^*\tilde{W})^\dagger c)^* v\right\} = \max_{v\in\mathcal{V}} \|v\|^2, \tag{1.80}$$

from which the theorem follows. □

Figure 1.8(b) shows a geometric interpretation of the minimax solution. The set $\mathcal{G}$ (solid segment) of feasible signals is an intersection of the ellipsoid defined

by $\|Lx\| \leq \rho$ and the subspace $S^*x = c$, which is orthogonal to $\mathcal{S}$. Clearly, for any reconstruction $\hat{x} \in \mathcal{G}$, the worst case signal x lies on the boundary of $\mathcal{G}$. Therefore, to minimize the worst case error, $\hat{x}$ must be the midpoint of the solid segment, as shown in the figure. The optimal reconstruction space $\tilde{\mathcal{W}}$ connects the recoveries corresponding to all possible sequences of samples c. This is equivalent to horizontally swapping the vertical dashed line in the figure and connecting the midpoints of the corresponding feasible sets $\mathcal{G}$.

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.

Example 1.3: Figure 1.9 compares the minimax (and LS) approach with bicubic interpolation in the context of image enlargement. The bicubic kernel of [35] is one of the most popular image re-sampling methods. In our experiment, a high-resolution image was down-sampled by a factor of 3. A continuous model was then fitted to the samples using both the minimax solution and the bicubic kernel of [35]. These models were re-sampled on a grid with 1/3 spacings to produce an image of the original size. The regularization operator was taken to be $L(\boldsymbol{\omega}) = \left((0.1\pi)^2 + \|\boldsymbol{\omega}\|^2\right)^{1.3}$, where $\boldsymbol{\omega}$ denotes the 2D frequency vector. In this example 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.

1.6.2 Constrained reconstruction

We next treat the problem of approximating x from its samples using a pre-specified set of reconstruction functions $\{w_n\}$. We will see that in this setup the LS and minimax recovery methods no longer coincide.

1.6.2.1 Least-squares approximation

In order to produce a solution $\hat{x} \in \mathcal{W}$, we modify the feasible set $\mathcal{G}$ of (1.62) to include only signals in $\mathcal{W}$:

$$\hat{x}_{\text{CLS}} = \arg\min_{x \in \tilde{\mathcal{G}}} \|S^*x - c\|^2, \tag{1.81}$$

where $\tilde{\mathcal{G}} = \{x : \|Lx\| \leq \rho, x \in \mathcal{W}\}$.

(a) Bicubic

(b) Minimax

Figure 1.9: Mandrill image rescaling: down-sampling by a factor of 3 using a rectangular sampling filter followed by upsampling back to the original dimensions using two interpolation methods. (a) The bicubic interpolation kernel leads to a blurry reconstruction with PSNR of 24.18dB. (b) The minimax method leads to a sharper reconstruction with PSNR of 24.39dB.

We have seen in Section 1.5 that without the constraint $\|Lx\| \leq \rho$, the set of solutions to (1.81) is given by

$$\mathcal{G} = \{x : x \in \mathcal{W}, S^*x = P_{\mathcal{R}(S^*W)}c\}. \tag{1.82}$$

We assume here that ρ is sufficiently large so that $\mathcal{G}$ contains at least one x. To choose one solution to (1.81) we minimize the smoothness measure $\|Lx\|$ over the set $\mathcal{G}$:

$$\hat{x}_{\text{CLS}} = \arg\min_{x \in \mathcal{G}} \|Lx\|. \tag{1.83}$$

Theorem 1.7. *The solution to problem (1.83) is given by*

$$\hat{x}_{CLS} = \hat{W}(S^*\hat{W})^\dagger c. \tag{1.84}$$

where now

$$\hat{W} = W(W^*L^*LW)^\dagger W^*S. \tag{1.85}$$

Proof. The proof of the theorem follows similar steps as in Section 1.6.1.1 and utilizes the fact that every signal in $\mathcal{G}$ is of the form $x = W((S^*W)^\dagger c + v)$, where $v \in \mathcal{N}(S^*W)$. □

Note that this solution is feasible, namely $\hat{x}_{\mathrm{CLS}} \in \mathcal{G}$:

$$\begin{aligned} S^*x_{\mathrm{CLS}} &= S^*W(W^*L^*LW)^\dagger W^*S(S^*W(W^*L^*LW)^\dagger W^*S)^\dagger c \\ &= P_{\mathcal{R}(S^*W(W^*L^*LW)^\dagger W^*S)}c \\ &= P_{\mathcal{R}(S^*W)}c. \end{aligned} \tag{1.86}$$

The last equality follows from the fact that $\mathcal{R}((W^*L^*LW)^\dagger) = \mathcal{N}(W^*L^*LW)^\perp = \mathcal{N}(LW)^\perp = \mathcal{N}(W)^\perp$ and similarly $\mathcal{N}((W^*L^*LW)^\dagger)^\perp = \mathcal{R}(W^*)$.

In contrast to subspace priors, here the constrained LS recovery does not generally relate to the unconstrained solution via an oblique projection. An exception is the case where $\mathcal{W} \oplus \mathcal{S}^\perp = \mathcal{H}$. As we have seen, in this case there exists a unique $x \in \mathcal{W}$ satisfying $S^*x = c$, which is equal to the oblique projection $E_{\mathcal{W}\mathcal{S}^\perp}x$. Since there is only one signal in the constraint set of problem (1.83), the smoothness measure in the objective does not play a role and the solution becomes $\hat{x}_{\mathrm{CLS}} = W(S^*W)^\dagger c$. In this setting, we can also use the fact that the unconstrained solution (1.65) satisfies $S^*x_{\mathrm{LS}} = c$, to write $\hat{x}_{\mathrm{CLS}} = W(S^*W)^\dagger S^*x_{\mathrm{LS}}$, recovering the relation we had in Section 1.5.2.1.

Another interesting scenario where L does not affect the solution is the case where $\mathcal{W}$ and $\mathcal{S}$ are SI spaces with generator $w(t)$ and $s(t)$ respectively and L is an LTI operator with frequency response $L(\omega)$. The operator $(W^*L^*LW)^\dagger$ then corresponds to the digital filter

$$\begin{cases} \frac{1}{\phi_{(LW)(LW)}(e^{j\omega})}, & \phi_{(LW)(LW)}(e^{j\omega}) \neq 0; \\ 0, & \phi_{(LW)(LW)}(e^{j\omega}) = 0 \end{cases} = \begin{cases} \frac{1}{\phi_{(LW)(LW)}(e^{j\omega})}, & \phi_{WW}(e^{j\omega}) \neq 0; \\ 0, & \phi_{WW}(e^{j\omega}) = 0, \end{cases} \tag{1.87}$$

where $\phi_{(LW)(LW)}(e^{j\omega})$ is given by (1.34) with $S(\omega)$ and $A(\omega)$ both replaced by $L(\omega)W(\omega)$ and $\phi_{WW}(e^{j\omega})$ is given by (1.34) with $S(\omega)$ and $A(\omega)$ both replaced by $W(\omega)$. The equality follows from the fact that L is assumed to satisfy (1.64) and thus $L(\omega)$ does not vanish anywhere. Therefore, it can be verified that $\hat{x}_{\mathrm{CLS}}(t)$ of (1.84) can be produced by filtering the sequence of samples $c[n]$ with

$$H(e^{j\omega}) = \begin{cases} \frac{1}{\phi_{SW}(e^{j\omega})}, & \phi_{SW}(e^{j\omega}) \neq 0, \phi_{WW}(e^{j\omega}) \neq 0; \\ 0, & \text{else}, \end{cases} \tag{1.88}$$

prior to reconstruction with $W(\omega)$. Here $\phi_{SW}(e^{j\omega})$ is given by (1.34) with $A(\omega)$ replaced by $W(\omega)$. It can be seen that (1.88) does not depend on $L(\omega)$, namely the smoothness prior does not affect the solution in the SI setting.

The situation where $\mathcal{W} \oplus \mathcal{S}^\perp = \mathcal{H}$, happens if and only if the supports of $S(\omega)$ and $W(\omega)$ are the same [32]. In this case it can be seen that (1.88) becomes

$$H(e^{j\omega}) = \begin{cases} \frac{1}{\phi_{SW}(e^{j\omega})}, & \phi_{SW}\left(e^{j\omega}\right) \neq 0; \\ 0, & \phi_{SW}\left(e^{j\omega}\right) = 0, \end{cases} \tag{1.89}$$

The resulting scheme is identical to the constrained LS reconstruction discussed in Section 1.6.2.1 in the context of subspace priors.

1.6.2.2 Minimax regret recovery

We next consider the extension of the minimax approach of Section 1.6.1.2 to the setup where $\hat{x}$ is constrained to lie in $\mathcal{W}$. Similar to the case of subspace priors, treated in Section 1.5.2.2, we consider here the minimization of the worst-case regret:

$$\hat{x}_{\mathrm{CMX}} = \arg\min_{\hat{x}\in\mathcal{W}} \max_{x\in\mathcal{G}} \left\|\hat{x} - P_{\mathcal{W}}x\right\|^2, \tag{1.90}$$

where $\mathcal{G}$ is given by (1.74).

Theorem 1.8. *The solution to (1.90) is given by*

$$\hat{x}_{CMX} = P_{\mathcal{W}}\tilde{W}(S^*\tilde{W})^\dagger c = P_{\mathcal{W}}\hat{x}_{MX}, \tag{1.91}$$

where $\tilde{W}$ is given by (1.66) and $\hat{x}_{MX} = \hat{x}_{LS}$ is the unconstrained solution given by (1.65).

Proof. The proof follows the exact same steps as in Section 1.6.1.2. □

This result is intuitive: When the output is constrained to the subspace $\mathcal{W}$, the minimax recovery is the orthogonal projection onto $\mathcal{W}$ of the minimax solution without the restriction. Recall that relation (1.91) is also true for subspace priors, as we have seen in Section 1.5.

Figure 1.10 shows a geometric interpretation of the minimax regret solution. As in the unconstrained scenario of Fig. 1.9, the feasible set of signals $\mathcal{G}$ is the vertical solid segment. Here, however, the reconstruction $\hat{x}$ is constrained to lie in the predefined space $\mathcal{W}$. The regret criterion (1.90) measures the deviation of $\hat{x}$ from $P_{\mathcal{W}}x$. The tilted solid segment is the projection of the feasible set $\mathcal{G}$ onto $\mathcal{W}$. For every reconstruction $\hat{x}$ in $\mathcal{W}$, the signal x leading to the worst regret corresponds to one of the endpoints of this set. Therefore, the minimal regret is attained if we choose $\hat{x}$ to be the midpoint of this segment. This solution is also the projection of the midpoint of $\mathcal{G}$ onto $\mathcal{W}$, *i.e.,* the projection of the unconstrained minimax reconstruction (1.65) onto $\mathcal{W}$.

When $\mathcal{S}$ and $\mathcal{W}$ are SI spaces and L is an LTI operator, the correction transform H corresponds to a digital filter $H(e^{j\omega})$. This filter can be determined by writing $H = (W^*W)^\dagger W^*\tilde{W}(S^*\tilde{W})^\dagger$, where $\tilde{W} = (L^*L)^{-1}S$ is the set transform corresponding to the unrestricted minimax solution. The operators W^*W, $W^*\tilde{W}$, and $S^*\tilde{W}$ correspond to the digital filters $\phi_{WW}(e^{j\omega})$, $\phi_{W\tilde{W}}(e^{j\omega})$ and $\phi_{S\tilde{W}}(e^{j\omega})$ respectively. The digital correction filter of Fig. 1.4 then becomes

$$H(e^{j\omega}) = \begin{cases} \frac{\phi_{W\tilde{W}}(e^{j\omega})}{\phi_{S\tilde{W}}(e^{j\omega})\phi_{WW}(e^{j\omega})}, & \phi_{S\tilde{W}}(e^{j\omega})\phi_{WW}(e^{j\omega}) \neq 0; \\ 0, & \text{else.} \end{cases} \tag{1.92}$$

In contrast to the constrained LS method, this filter depends on $L(\omega)$ so that the prior does affect the solution. The next example demonstrates the effectiveness of this filter in an image processing task.

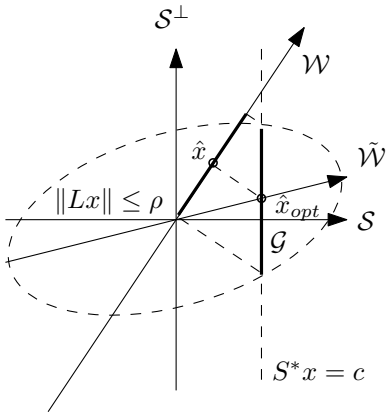

Figure 1.10: Geometric interpretation of minimax regret recovery in a predefined reconstruction space $\mathcal{W}$.

Example 1.4: In Fig. 1.11 we demonstrate the difference between the LS and minimax-regret methods in an image enlargement task. The setup is the same as that of Fig. 1.9 only now the reconstruction filter is constrained to be a triangular kernel corresponding to linear interpolation. With this interpolation kernel, the direct-sum condition $L_2 = \mathcal{W} \oplus \mathcal{S}^\perp$ is satisfied. It can be seen that the error of the minimax regret recovery is only 0.7dB less than the unconstrained minimax shown in Fig. 1.9. The constrained LS 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 (1.92) can be best understood by examining the case where our only prior on the signal is that it is norm-bounded, that is, when L is the identity operator I. This scenario was thoroughly investigated in [32]. Setting $L(\omega) = 1$ in (1.92), the correction filter becomes

$$H(e^{j\omega}) = \begin{cases} \frac{\phi_{WS}(e^{j\omega})}{\phi_{SS}(e^{j\omega})\phi_{WW}(e^{j\omega})}, & \phi_{SS}(e^{j\omega})\phi_{WW}(e^{j\omega}) \neq 0; \\ 0, & \text{else}, \end{cases} \tag{1.93}$$

since from (1.72), $\tilde{w}(t) = s(t)$. Applying the Cauchy-Schwartz inequality to the numerator of (1.93) and to the denominator of (1.88), it is easy to see that the magnitude of the minimax regret filter (1.93) is smaller than that of the constrained LS filter (1.88) at all frequencies. This property renders the minimax regret approach more resistant to noise in the samples $c[n]$, since perturbations

(a) Consistent (b) Minimax regret

Figure 1.11: Mandrill image rescaling: down-sampling by a factor of 3 using a rectangular sampling filter followed by upsampling back to the original dimensions using the LS and minimax regret methods. (a) The LS approach over-enhances the high frequencies and results in a PSNR of 22.51dB. (b) The minimax regret method leads to a smoother reconstruction with PSNR of 23.69dB.

in $\hat{x}(t)$ caused by errors in $c[n]$ are always smaller in the minimax regret method than in the consistent approach.

In Fig. 1.12 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 $P_{\mathcal{S}}x$. 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.

When x is known to lie in $\mathcal{S}$, it follows from the previous section that the minimal error can be obtained by using (1.60) with $A = S$. The resulting filter coincides with the minimax regret filter of (1.89), implying that the regret approach minimizes the squared-error over all $x \in \mathcal{S}$.

In [32] tight bounds on the error resulting from the constrained LS and minimax regret methods are developed for the case where $\mathcal{H} = \mathcal{W} \oplus \mathcal{S}^{\perp}$. We omit the technical details here and only summarize the main conclusions. We first recall that if we know a priori that x lies in a subspace $\mathcal{A}$ such that $\mathcal{H} = \mathcal{A} \oplus \mathcal{S}^{\perp}$, then the filter (1.60) will yield the minimal error approximation of x and therefore is optimal in the squared-norm sense. When $\mathcal{A} = \mathcal{S}$ this strategy reduces to the minimax regret method, while if $\mathcal{A} = \mathcal{W}$, then we obtain the constrained LS reconstruction.

When no prior subspace knowledge is given, the regret approach is preferable if the spaces $\mathcal{S}$ and $\mathcal{W}$ are sufficiently far apart, or if x has enough energy in $\mathcal{S}$. These results are intuitive as illustrated geometrically in Fig. 1.12. In Fig. 1.12(a) we depict the constrained LS and regret reconstruction when $\mathcal{W}$ is far from $\mathcal{S}$.

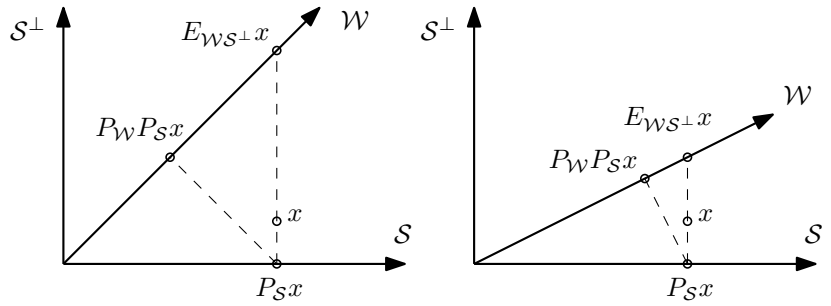

Figure 1.12: Comparison of minimax regret reconstruction and constrained LS reconstruction for two different choices of $\mathcal{W}$ satisfying $\mathcal{H} = \mathcal{W} \oplus \mathcal{S}^\perp$. (a) The minimax strategy ($P_\mathcal{W}P_\mathcal{S}x$) is preferable to LS ($E_{\mathcal{W}\mathcal{S}^\perp}x$) when $\mathcal{W}$ is 'far' from $\mathcal{S}$. (b) Both methods lead to errors on the same order of magnitude when $\mathcal{W}$ is 'close' to $\mathcal{S}$.

Table 1.2: Reconstruction from noiseless samples

	Unconstrained ($\hat{x} \in \mathcal{H}$)		Constrained ($\hat{x} \in \mathcal{W}$)	
Prior	Least-Squares	Minimax	Least-Squares	Minimax
$x \in \mathcal{A}$	$\hat{x}^{\rm sub}_{\rm LS} = A(S^*A)^\dagger c$	$\hat{x}^{\rm sub}_{\rm MX} = \hat{x}^{\rm sub}_{\rm LS}$	$\hat{x}^{\rm sub}_{\rm CLS} = W(S^*W)^\dagger S^* \hat{x}^{\rm sub}_{\rm LS}$	$\hat{x}^{\rm sub}_{\rm CMX} = P_\mathcal{W}\hat{x}^{\rm sub}_{\rm MX}$
$\|Lx\| \le \rho$	$\hat{x}^{\rm smo}_{\rm LS} = \tilde{W}(S^*\tilde{W})^\dagger c$	$\hat{x}^{\rm smo}_{\rm MX} = \hat{x}^{\rm smo}_{\rm LS}$	$\hat{x}^{\rm smo}_{\rm CLS} = \hat{W}(S^*\hat{W})^\dagger S^* \hat{x}^{\rm smo}_{\rm LS}$	$\hat{x}^{\rm smo}_{\rm CMX} = P_\mathcal{W}\hat{x}^{\rm smo}_{\rm MX}$

As can be seen in the figure, in this case the error resulting from the LS solution is large with respect to the regret approximation error. In Fig. 1.12(b), $\mathcal{W}$ and $\mathcal{S}$ are close, and the errors have roughly the same magnitude.

1.7 Comparison of the various scenarios

Table 1.2 summarizes the reconstruction techniques developed in Sections 1.5 and 1.6. We use the superscripts 'sub' and 'smo' to signify whether a solution corresponds to a subspace or a smoothness prior. The transformations $\tilde{W}$ and $\hat{W}$ are given by

$$\tilde{W} = (L^*L)^{-1}S$$

and

$$\hat{W} = W(W^*L^*LW)^\dagger W^*S$$

respectively.

This table highlights the key observations discussed in the previous sections. We begin by examining the case in which no constraint is imposed on $\hat{x}$, shown

Table 1.3: Reconstruction from noiseless samples under direct sum assumptions

	Unconstrained ($\hat{x} \in \mathcal{H}$)		Constrained ($\hat{x} \in \mathcal{W}$)	
Prior	Least-Squares	Minimax	Least-Squares	Minimax
$x \in \mathcal{A}$	$\hat{x}_{\rm LS}^{\rm sub} = x$	$\hat{x}_{\rm MX}^{\rm sub} = \hat{x}_{\rm LS}^{\rm sub}$	$\hat{x}_{\rm CLS}^{\rm sub} = E_{\mathcal{W}\mathcal{S}^\perp}\hat{x}_{\rm LS}^{\rm sub}$	$\hat{x}_{\rm CMX}^{\rm sub} = P_{\mathcal{W}}\hat{x}_{\rm MX}^{\rm sub}$
$\|Lx\| \leq \rho$	$\hat{x}_{\rm LS}^{\rm smo} = E_{\tilde{\mathcal{W}}\mathcal{S}^\perp}x$	$\hat{x}_{\rm MX}^{\rm smo} = \hat{x}_{\rm LS}^{\rm smo}$	$\hat{x}_{\rm CLS}^{\rm smo} = E_{\mathcal{W}\mathcal{S}^\perp}\hat{x}_{\rm LS}^{\rm smo}$	$\hat{x}_{\rm CMX}^{\rm smo} = P_{\mathcal{W}}\hat{x}_{\rm MX}^{\rm smo}$

in columns 1 and 2. First, we see that in this situation both the LS and minimax reconstructions coincide. This property holds true both for subspace and smoothness priors. In Section 1.8 we show that this is not the case when the samples are noisy. Second, smoothness-prior recovery (row 2) has the same structure as subspace-prior recovery (row 1) with $\tilde{W}$ replacing A. Therefore, we can interpret $\tilde{\mathcal{W}} = \mathcal{R}(\tilde{W})$ as the optimal reconstruction space associated with the smoothness prior. Finally, we note that for certain subspace priors, perfect recovery can be achieved, leading to $\hat{x}_{\rm LS}^{\rm sub} = x$. Specifically, this happens if the sampling space $\mathcal{S}$ and the prior space $\mathcal{A}$ satisfy the direct sum condition $\mathcal{H} = \mathcal{A} \oplus \mathcal{S}^\perp$. In the smoothness prior case, the direct sum condition $\mathcal{H} = \tilde{\mathcal{W}} \oplus \mathcal{S}^\perp$ does not imply perfect recovery because the original x does not necessarily lie in $\tilde{\mathcal{W}}$. When the direct sum holds, however, recovery can be interpreted as an oblique projection of the (unknown) signal x onto the optimal reconstruction space $\tilde{\mathcal{W}}$, namely $\hat{x}_{\rm LS}^{\rm smo} = E_{\tilde{\mathcal{W}}\tilde{\mathcal{S}}^\perp}x$.

We now examine the case in which the recovery is constrained to lie in $\mathcal{W}$ (columns 3 and 4). These solutions are expressed in Table 1.2 in terms of the unconstrained reconstructions (columns 1 and 2). Here the minimax regret solutions (column 4) are related to the unconstrained recoveries (column 2) via an orthogonal projection onto the reconstruction space $\mathcal{W}$. This implies that $\|\hat{x}_{\rm CMX}\| \leq \|\hat{x}_{\rm MX}\|$. The constrained LS solutions, on the other hand, possess a different structure. When the sampling and reconstruction spaces satisfy the direct-sum $\mathcal{H} = \mathcal{W} \oplus \mathcal{S}^\perp$, both constrained LS solutions of column 3 become $\hat{x}_{\rm CLS} = E_{\mathcal{W}\mathcal{S}^\perp}\hat{x}_{\rm LS}$. This has several implications. First, in contrast to an orthogonal projection, an oblique projection may lead to solutions with arbitrary large norm, given that $\mathcal{W}$ is sufficiently far apart from $\mathcal{S}$. Therefore, the error in the constrained LS framework is not guaranteed to be bounded, unless a bound on the 'distance' between $\mathcal{S}$ and $\mathcal{W}$ is known a-priori. Second, this implies that the recovery does not depend on the prior, *i.e.,* $\hat{x}_{\rm CLS}^{\rm sub}$ is not a function of A and $\hat{x}_{\rm CLS}^{\rm smo}$ does not depend on L. These properties are clearly undesirable and can lead to unsatisfactory results in practical applications, as demonstrated in Example 1.4.

Table 1.3 summarizes the recovery formulae obtained under the direct sum assumptions discussed above. The expressions in column 1, rows 1 and 2, are true when $\mathcal{H} = \mathcal{A} \oplus \mathcal{S}^\perp$ and $\mathcal{H} = \tilde{\mathcal{W}} \oplus \mathcal{S}^\perp$ respectively. The recoveries of column 3 are obtained under the assumption that $\mathcal{H} = \mathcal{W} \oplus \mathcal{S}^\perp$.

Finally, we note that all the recovery techniques discussed thus far can be easily implemented in SI spaces. Specifically, suppose that $\mathcal{S}$, $\mathcal{A}$ and $\mathcal{W}$ are SI spaces with generators $s(t)$, $a(t)$ and $w(t)$ respectively. Moreover, assume that L is an

Table 1.4: Reconstruction from noiseless samples in SI spaces

	Unconstrained ($\hat{x} \in \mathcal{H}$)		Constrained ($\hat{x} \in \mathcal{W}$)	
Prior	Least-Squares	Minimax	Least-Squares	Minimax
$x \in \mathcal{A}$	(1.33)	(1.33)	(1.50)	(1.60)
$\|Lx\| \leq \rho$	(1.73), (1.72)	(1.73), (1.72)	(1.88)	(1.92)

LTI operator corresponding to the filter $L(\omega)$. Then all the reconstruction methods of Table 1.2 can be implemented by digitally filtering the samples $c[n]$ prior to reconstruction, as depicted in Fig. 1.4. The resulting interpolation methods are summarized in Table 1.4. The numbers in the table indicate the equation numbers containing the reconstruction formulae of the digital correction filter and reconstruction kernel. The optimal kernel corresponding to unconstrained recovery with a subspace prior (row 1, columns 1 and 2) is $a(t)$, while the interpolation kernel in the constrained case (columns 3 and 4) is $w(t)$.

The direct sum conditions, under which Table 1.3 was constructed, can be easily verified in SI spaces, as explained in Section 1.5.1.1. Specifically, for two SI spaces $\mathcal{A}$ and $\mathcal{S}$, the condition $\mathcal{H} = \mathcal{A} \oplus \mathcal{S}^{\perp}$ is satisfied if and only if [46] the supports $\mathcal{I}_{\mathcal{A}}$ and $\mathcal{I}_{\mathcal{S}}$ of $\phi_{SS}(e^{j\omega})$ and $\phi_{AA}(e^{j\omega})$ respectively coincide, and there exists a constant $\alpha > 0$ such that $|\phi_{SA}(e^{j\omega})| > \alpha$, for all ω in $\mathcal{I}_{\mathcal{A}}$. The filters $\phi_{SS}(e^{j\omega})$, $\phi_{AA}(e^{j\omega})$, and $\phi_{SA}(e^{j\omega})$ are defined in (1.34).

1.8 Sampling with noise

We now extend the approaches of the previous sections to the case in which the samples are perturbed by noise. Specifically we assume that $c = S^*x + u$, where $u[n]$ is an unknown noise sequence.

To approach the noisy setup within the LS framework, we need to minimize the error-in-samples $\|S^*x - c\|^2$ over the set of feasible signals. Thus, the optimization problems (1.24), (1.45), (1.62) and (1.81), which correspond to the unconstrained and constrained subspace and smoothness scenarios, remain valid here too. However, note that to solve these problems we assumed that signals x for which $S^*x = c$ (or $S^*x = P_{\mathcal{R}(W^*S)}c$ in the constrained setting) are included in the feasible set. When the samples are noisy, this is not necessarily true so that, for example, the optimal value of the unconstrained problem $\min_{x\in\mathcal{A}} \|S^*x - c\|^2$ is no longer 0. Nevertheless, it can be easily shown that the solutions we obtained under the subspace prior assumption (problems (1.24) and (1.45)) and in the constrained smoothness setting (problem (1.81)) remain the same. Furthermore, it can be shown that in the unconstrained smoothness scenario (problem (1.62)), this fact does not change the optimal reconstruction space $\tilde{\mathcal{W}}$ of (1.66), but only the expansion coefficients of $\hat{x}$ in $\tilde{\mathcal{W}}$. Interestingly, this property holds even when

the ℓ_2-norm in the error-in-samples term $\|S^*x - c\|$ is replaced by an ℓ_p-norm with arbitrary $p \in [1, \infty]$ [11].

Therefore, we focus our attention on the extension of the *minimax* recovery techniques to the noisy case. To keep the exposition simple, we will thoroughly examine only the smoothness prior scenarios of Section 1.6. The subspace prior problems of Section 1.5 can be treated in the exact same manner. We thus assume in the sequel that $\|Lx\| \leq \rho$ for some $\rho \geq 0$. In this setting the solution no longer lies in the reconstruction space $\tilde{\mathcal{W}}$ of (1.66). Moreover, the resulting problems generally do not admit a closed-form solution and must be solved using numerical optimization methods. Consequently, we will narrow the discussion from signals in an arbitrary Hilbert space $\mathcal{H}$ to signals lying in $\mathbb{R}^n$ or $\mathbb{C}^n$.

Assume that the samples c are noisy so that our only information is that $\|S^*x - c\| \leq \alpha$ for some value of α. The extension of the minimax problem (1.75) in the unconstrained scenario to the noisy case is

$$\hat{x}_{\mathrm{MX}} = \arg\min_{\hat{x}} \max_{x \in \mathcal{G}} \|\hat{x} - x\|^2, \tag{1.94}$$

where

$$\mathcal{G} = \{x : \|S^*x - c\| \leq \alpha, \|Lx\| \leq \rho\}. \tag{1.95}$$

Similarly, the counterpart of (1.90), where the solution is constrained to lie in $\mathcal{W}$, is $\hat{x}_{\mathrm{CMX}} = W\hat{d}_{\mathrm{CMX}}$, with

$$\hat{d}_{\mathrm{CMX}} = \arg\min_{d} \max_{x \in \mathcal{G}} \|Wd - P_{\mathcal{W}}x\|^2 . \tag{1.96}$$

To solve (1.94) and (1.96) we replace the inner maximization by its dual function. This will result in a minimization problem that can be combined with the outer minimization. In order to follow this method, we need to be able to establish strong duality of the inner maximization with its dual function. The maximization in (1.94) and (1.96) is a special case of a nonconvex quadratic optimization problem with two quadratic constraints. The nonconvexity is the result of the fact that we are maximizing a convex quadratic function, rather than minimizing it. Nonconvex quadratic optimization problems have been studied extensively in the optimization literature. Below, we first briefly survey some of the main results on quadratic optimization, relevant to our problem. We then show how they can be used to develop a robust recovery method.

1.8.1 Quadratic optimization problems

This simplest class of quadratic optimization problems is the minimization of a single (possibly nonconvex) quadratic function subject to one quadratic constraint. A special well-studied case is that of the trust region algorithm for unconstrained optimization, which has the form [50, 51, 52, 53, 54, 55]:

$$\min_{x \in \mathbb{R}^n} \{x^T Bx + 2g^T x : \|x\|^2 \leq \delta\}, \tag{1.97}$$

where B is not necessarily nonnegative definite so that the problem is not generally convex. The dual of (1.97) is given by the semidefinite program (SDP) [56]

$$\max_{\alpha,\lambda}\left\{\lambda \ : \ \begin{pmatrix} B+\alpha I & g \\ g^T & -\alpha\delta-\lambda \end{pmatrix} \succeq 0, \alpha \geq 0\right\}. \tag{1.98}$$

Problem (1.97) enjoys many useful and attractive properties. It is known that it admits no duality gap and that its semidefinite relaxation (SDR) is tight. Moreover, there exist a set of necessary and sufficient conditions that guarantee optimality of a solution to (1.97), which can be extracted from the dual solution $\bar{\alpha}$. These results all extend to the case in which the problem is to optimize an arbitrary quadratic function subject to a single quadratic constraint, where both quadratic forms are not necessarily convex.

Unfortunately, in general these results cannot be generalized to the case in which the constraint set consists of two quadratic restrictions. More specifically, consider the following quadratic problems:

$$(Q2P_{\mathbb{C}}) \quad \min_{z\in\mathbb{C}^n}\{f_3(z) : f_1(z) \geq 0, f_2(z) \geq 0\}, \tag{1.99}$$

$$(Q2P_{\mathbb{R}}) \quad \min_{x\in\mathbb{R}^n}\{f_3(x) : f_1(x) \geq 0, f_2(x) \geq 0\}. \tag{1.100}$$

In the real case each function $f_j : \mathbb{R}^n \to \mathbb{R}$ is defined by $f_j(x) = x^T A_j x + 2b_j^T x + c_j$ with $A_j = A_j^T \in \mathbb{R}^{n\times n}, b_j \in \mathbb{R}^n$ and $c_j \in \mathbb{R}$. In the complex setting, $f_j : \mathbb{C}^n \to \mathbb{R}$ is given by $f_j(z) = z^* A_j z + 2\Re(b_j^* z) + c_j$, where A_j are Hermitian matrices, *i.e.,* $A_j = A_j^*$, $b_j \in \mathbb{C}^n$ and $c_j \in \mathbb{R}$. We distinguish between the real and complex cases since we will see that different strong duality results apply in both settings. In particular, there are stronger results for complex quadratic problems than for their real counterparts.

The problem $(Q2P_{\mathbb{R}})$ appears as a subproblem in some trust region algorithms for constrained optimization [57, 58, 59, 60, 61] where the original problem is to minimize a general nonlinear function subject to equality constraints. Unfortunately, in general the strong duality results in the case of a single constraint cannot be extended to the case of two quadratic restrictions $(Q2P_{\mathbb{R}})$. Indeed, it is known that the SDR of $(Q2P_{\mathbb{R}})$ is not necessarily tight [62, 60]. An exception is the case in which the functions f_1, f_2 and f_3 are all homogenous quadratic functions and there exists a positive definite linear combination of the matrices A_j [62]. Another setting in which strong duality is guaranteed is derived in [44] and will be discussed below.

Quadratic optimization in the complex domain is simpler. In [44] it is shown that under some mild conditions strong duality holds for the complex valued problem $(Q2P_{\mathbb{C}})$ and that its semidefinite relaxation is tight. This result is based on the extended version of the S-lemma derived by Fradkov and Yakubovich [63]. The standard Lagrangian dual of $(Q2P_{\mathbb{C}})$ is given by

$$(D_{\mathbb{C}}) \quad \max_{\alpha\geq 0,\beta\geq 0,\lambda}\left\{\lambda \ \middle| \ \begin{pmatrix} A_3 & b_3 \\ b_3^* & c_3-\lambda \end{pmatrix} \succeq \alpha\begin{pmatrix} A_1 & b_1 \\ b_1^* & c_1 \end{pmatrix} + \beta\begin{pmatrix} A_2 & b_2 \\ b_2^* & c_2 \end{pmatrix}\right\}. \tag{1.101}$$

Problem $(D_{\mathbb{C}})$ is sometimes called Shor's relaxation [64]. Theorem 1.9 below states that if problem $(Q2P_{\mathbb{C}})$ is strictly feasible then $\mathrm{val}(Q2P_{\mathbb{C}}) = \mathrm{val}(D_{\mathbb{C}})$ *even* in the case where the value is equal to $-\infty$.

Theorem 1.9. *[44] Suppose that problem $(Q2P_{\mathbb{C}})$ is strictly feasible,* i.e., *there exists $\tilde{z} \in \mathbb{C}^n$ such that $f_1(\tilde{z}) > 0, f_2(\tilde{z}) > 0$. Then,*

1. *if $val(Q2P_{\mathbb{C}})$ is finite then the maximum of problem $(D_{\mathbb{C}})$ is attained and $val(Q2P_{\mathbb{C}}) = val(D_{\mathbb{C}})$.*
2. *$val(Q2P_{\mathbb{C}}) = -\infty$ if and only if $(D_{\mathbb{C}})$ is not feasible.*

Necessary and sufficient optimality conditions similar to those known for (1.97) where also derived in [44]. These conditions can be used to calculate the optimal solution of $(Q2P_{\mathbb{C}})$ from the dual solution.

It is interesting to note that the dual problem to $(D_{\mathbb{C}})$ is the so-called SDR of $(Q2P_{\mathbb{C}})$:

$$(SDR_{\mathbb{C}}) \quad \min_{Z}\{\mathrm{Tr}(ZM_3) : \mathrm{Tr}(ZM_1) \geq 0, \mathrm{Tr}(ZM_2) \geq 0, Z_{n+1,n+1} = 1, Z \succeq 0\}, \tag{1.102}$$

where

$$M_j = \begin{pmatrix} A_j & b_j \\ b_j^* & c_j \end{pmatrix}. \tag{1.103}$$

If both problems $(Q2P_{\mathbb{C}})$ and $(D_{\mathbb{C}})$ are strictly feasible, then problems $(Q2P_{\mathbb{C}})$,$(D_{\mathbb{C}})$ and $(SDR_{\mathbb{C}})$ (problems (1.99),(1.101) and (1.102) respectively) attain their solutions and

$$\mathrm{val}(Q2P_{\mathbb{C}}) = \mathrm{val}(D_{\mathbb{C}}) = \mathrm{val}(SDR_{\mathbb{C}}). \tag{1.104}$$

The real valued problem $(Q2P_{\mathbb{R}})$ is more difficult to handle. In contrast to the complex case, strong duality results are, generally speaking, not true for $(Q2P_{\mathbb{R}})$. It is not known whether $(Q2P_{\mathbb{R}})$ is a tractable problem or not and in that respect, if there is an efficient algorithm for finding its solution. If the constraints of $(Q2P_{\mathbb{R}})$ are convex then the complex valued problem $(Q2P_{\mathbb{C}})$, considered as a relaxation of $(Q2P_{\mathbb{R}})$, can produce an approximate solution. Although strong duality results do not hold generally in the real case, a sufficient condition can be developed to ensure zero duality gap (and tightness of the semidefinite relaxation) for $(Q2P_{\mathbb{R}})$ [44]. This result is based on the connection between the image of the real and complex spaces under a quadratic mapping, and is given in terms of the dual optimal values.

The dual problem to $(QP_{\mathbb{R}})$ is

$$(D_{\mathbb{R}}) \quad \max_{\alpha\geq 0,\beta\geq 0,\lambda} \left\{ \lambda \;\middle|\; \begin{pmatrix} A_3 & b_3 \\ b_3^T & c_3 - \lambda \end{pmatrix} \succeq \alpha \begin{pmatrix} A_1 & b_1 \\ b_1^T & c_1 \end{pmatrix} + \beta \begin{pmatrix} A_2 & b_2 \\ b_2^T & c_2 \end{pmatrix} \right\}. \tag{1.105}$$

Note that this is exactly the same as problem $(D_{\mathbb{C}})$ in (1.101), where here we used the fact that the data is real and therefore $b_j^* = b_j^T$. The SDR in this case

is given by

$$(SDR_{\mathbb{R}}) \quad \min_{X}\{\mathrm{Tr}(XM_3) : \mathrm{Tr}(XM_1) \geq 0, \mathrm{Tr}(XM_2) \geq 0, X_{n+1,n+1} = 1, X \succeq 0\}. \tag{1.106}$$

Suppose that both problems $(Q2P_{\mathbb{R}})$ and $(D_{\mathbb{R}})$ are strictly feasible and there exists real values $\hat{\alpha}, \hat{\beta}$ such that

$$\hat{\alpha}A_1 + \hat{\beta}A_2 \succ 0. \tag{1.107}$$

Let $(\bar{\lambda}, \bar{\alpha}, \bar{\beta})$ be an optimal solution of the dual problem $(D_{\mathbb{R}})$. If

$$\dim\left(\mathcal{N}(A_3 - \bar{\alpha}A_1 - \bar{\beta}A_2)\right) \neq 1 \tag{1.108}$$

then $\mathrm{val}(Q2P_{\mathbb{R}}) = \mathrm{val}(D_{\mathbb{R}}) = \mathrm{val}(SDR_{\mathbb{R}})$ and there exists a real valued solution to the complex valued problem $(Q2P_{\mathbb{C}})$.

1.8.2 Minimax recovery using SDP relaxation

We now show how the strong duality results developed in the previous section can be used to solve the recovery problems (1.94) and (1.96). Our general approach is to replace the inner maximization by its dual [65]. Over the complex domain, strong duality holds, and the resulting problems are exact representations of (1.94) and (1.96). Over the reals, this leads to an approximation, however, in practice it is pretty tight and yields good recovery results. Alternatively, we can use an SDR approach to replace the inner maximization by its relaxation. The resulting problem is a convex-concave saddle point program which can be further simplified by relying on Proposition 1.1 [66]. Both derivations are equivalent since the dual problem of the inner maximization is also the dual of the (convex) SDR [56]. Here we follow the relaxation approach since its derivation is simpler.

Instead of focusing on our particular problem, we treat a general minimax formulation with a quadratic objective, and two quadratic constraints:

$$\min_{\hat{x}} \max_{x}\{\|A\hat{x} - Qx\|^2 : f_i(x) \leq 0, 1 \leq i \leq 2\}, \tag{1.109}$$

where

$$f_i(x) \stackrel{\triangle}{=} x^* A_i x + 2\Re\{b_i^* x\} + c_i. \tag{1.110}$$

Clearly (1.94) and (1.96) are special cases of (1.109) with

$$A_1 = SS^*, b_1 = -Sc, c_1 = \|c\|^2 - \alpha^2, A_2 = L^*L, b_2 = 0, c_2 = -\rho^2. \tag{1.111}$$

The difference between the two problems is in the matrices A and Q. In (1.94) we have $A = Q = I$, whereas in (1.96) $A = W$ and $Q = P_{\mathcal{W}} = W(W^*W)^{\dagger}W^*$.

In order to develop a solution to (1.109) we first consider the inner maximization:

$$\max_{x}\{\|A\hat{x} - Qx\|^2 : f_i(x) \leq 0, 1 \leq i \leq 2\}, \tag{1.112}$$

which is a special case of quadratic optimization with 2 quadratic constraints. Denoting $\Delta = xx^*$, (1.112) can be written equivalently as

$$\max_{(\Delta,x)\in\mathcal{G}} \{\|A\hat{x}\|^2 - 2\Re\{\hat{x}^* A^* Q x\} + \mathrm{Tr}(Q^* Q\Delta)\}, \tag{1.113}$$

where

$$\mathcal{G} = \{(\Delta, x) : f_i(\Delta, x) \le 0,\ 1 \le i \le 2, \Delta = xx^*\}, \tag{1.114}$$

and we defined

$$f_i(\Delta, x) = \mathrm{Tr}(A_i\Delta) + 2\Re\{b_i^* x\} + c_i, \quad 1 \le i \le 2. \tag{1.115}$$

The objective in (1.113) is concave (linear) in (Δ, x), but the set $\mathcal{G}$ is not convex. To obtain a relaxation of (1.113) we may replace $\mathcal{G}$ by the convex set

$$\mathcal{T} = \{(\Delta, x) : f_i(\Delta, x) \le 0,\ 1 \le i \le 2, \Delta \succeq xx^*\}. \tag{1.116}$$

Indeed, using Schur's lemma [67, p. 28] $\Delta \succeq xx^*$ can be a written as a linear matrix inequality. Our relaxation of (1.109) is the solution to the resulting minimax problem:

$$\min_{\hat{x}} \max_{(\Delta,x)\in\mathcal{T}} \{\|A\hat{x}\|^2 - 2\Re\{\hat{x}^* A^* Q x\} + \mathrm{Tr}(Q^* Q\Delta)\}. \tag{1.117}$$

The objective in (1.117) is concave (linear) in Δ and x and convex in $\hat{x}$. Furthermore, the set $\mathcal{T}$ is bounded. Therefore, from Proposition 1.1 we can replace the order of the minimization and maximization, resulting in the equivalent problem

$$\max_{(\Delta,x)\in\mathcal{T}} \min_{\hat{x}} \{\|A\hat{x}\|^2 - 2\Re\{\hat{x}^* A^* Q x\} + \mathrm{Tr}(Q^* Q\Delta)\}. \tag{1.118}$$

The inner minimization is a simple quadratic problem. Expressing the objective as $\|A\hat{x} - Qx\|^2$, it can be seen that its solution satisfies $A\hat{x} = P_{\mathcal{A}} Q x$, where $\mathcal{A} = \mathcal{R}(A)$. Substituting this result into (1.118), our problem reduces to

$$\max_{(\Delta,x)\in\mathcal{T}} \{-\|P_{\mathcal{A}} Q x\|^2 + \mathrm{Tr}(Q^* Q\Delta)\}, \tag{1.119}$$

Problem (1.119) is a convex optimization problem with a concave objective and linear matrix inequality constraints and can therefore be solved easily using standard software packages. The approximate minimax solution to (1.109) is the x-part of the solution to (1.119). When (1.109) is defined over the complex domain, then this solution is exact. In the real case, it will be exact when condition (1.108) is satisfied.

Instead of solving (1.119) we may consider its dual function. Since (1.119) is convex, strong duality holds. For simplicity, we will assume that $L^* L$ is invertible so that $A_2 \succ 0$ in our case. We will also use the fact that in our setting, $Q = P_{\mathcal{A}}$, where $\mathcal{A} = \mathcal{R}(A)$. This leads to a simple explicit expression for the solution $\hat{x}$:

Theorem 1.10. *Assume that at least one of the matrices $\{A_i\}$ is strictly positive definite and that $Q = P_{\mathcal{A}}$. Then the solution to (1.119) is given by*

$$\hat{x} = -\left(\sum_{i=1}^{2} \alpha_i A_i\right)^{-1} \left(\sum_{i=1}^{2} \alpha_i b_i\right), \tag{1.120}$$

where (α_1, α_2) is an optimal solution of the following convex optimization problem in 2 variables:

$$\begin{aligned} \min_{\alpha_i} & \left\{ \sum_{i=1}^{2} \alpha_i b_i^* \left(\sum_{i=1}^{2} \alpha_i A_i\right)^{-1} \sum_{i=1}^{2} \alpha_i b_i - \sum_{i=1}^{2} c_i \alpha_i \right\} \\ s.t. \quad & \sum_{i=1}^{2} \alpha_i A_i \succeq Q^* Q, \\ & \alpha_i \geq 0, \quad 1 \leq i \leq 2. \end{aligned} \tag{1.121}$$

Proof. To prove the theorem we show that (1.121) is the dual of (1.119). Since (1.119) is convex and strictly feasible, its optimal value is equal to that of its dual problem. To compute the dual, we first form the Lagrangian:

$$\begin{aligned} \mathcal{L} = & -\|P_{\mathcal{A}} Q x\|^2 + \mathrm{Tr}(Q^* Q \Delta) + \mathrm{Tr}(\Pi(\Delta - x x^T)) \\ & - \sum_{i=1}^{2} \alpha_i (\mathrm{Tr}(A_i \Delta) + 2\Re\{b_i^* x\} + c_i), \end{aligned} \tag{1.122}$$

where $\alpha_i \geq 0$ and $\Pi \succeq 0$ are the dual variables. The maximization of $\mathcal{L}$ with respect to x yields

$$x = -(Q^* P_{\mathcal{A}} Q + \Pi)^{-1} \sum_{i=1}^{2} \alpha_i b_i. \tag{1.123}$$

The derivative with respect to Δ yields,

$$Q^* Q + \Pi = \sum_{i=1}^{2} \alpha_i A_i. \tag{1.124}$$

Using the fact that $Q^* P_{\mathcal{A}} Q = Q^* Q$ and that $Q^*(P_{\mathcal{A}} - I)Q = 0$, the combination of (1.123) and (1.124) yields (1.120).

Next we note that since $\Pi \succeq 0$, we must have from (1.124) that $\sum_{i=0}^{k} \alpha_i A_i \succeq Q^* Q$. Finally, substituting (1.123) and (1.124) into (1.122), we obtain the dual problem (1.121). □

Returning to our reconstruction problems, the substitution of (1.111) in Theorem 1.10 implies that $\hat{x}_{\mathrm{MX}}$ of (1.94) is given by

$$\hat{x}_{\mathrm{MX}} = (SS^* + \lambda L^* L)^{-1} Sc, \tag{1.125}$$

where $\lambda = \alpha_2/\alpha_1$. If $\alpha_1 = 0$ then $\hat{x}_{\mathrm{MX}} = 0$. Similarly, $\hat{x}_{\mathrm{CMX}}$ corresponding to (1.96) is given by

$$\hat{x}_{\mathrm{CMX}} = P_{\mathcal{W}} \hat{x}_{\mathrm{MX}}, \tag{1.126}$$

where we used the fact that $A = W$ and $Q = P_{\mathcal{W}}$ in this case. This shows that, as in the noiseless scenarios of Sections 1.5 and 1.6, here too the constrained minimax regret solution relates to the unconstrained minimax recovery via an orthogonal projection.

Problem (1.121) can be cast as an SDP:

$$\begin{array}{ll} \min_{\alpha_i} & \left\{ t - \sum_{i=1}^{2} c_i \alpha_i \right\} \\ \text{s.t.} & \begin{pmatrix} Q^*(P_{\mathcal{A}} - I)Q + \sum_{i=1}^{2} \alpha_i A_i & \sum_{i=1}^{2} \alpha_i b_i \\ \sum_{i=1}^{2} \alpha_i b_i^* & t \end{pmatrix} \succeq 0, \\ & \sum_{i=1}^{2} \alpha_i A_i \succeq Q^* Q, \\ & \alpha_i \geq 0, \quad 1 \leq i \leq 2. \end{array} \tag{1.127}$$

This SDP can be solved by one of the many available SDP solvers such as the Self-Dual-Minimization (SeDuMi) package [68] or CVX [69].

It can be shown that both the true minimax solution and its approximation are feasible, namely they satisfy the quadratic constraints [66]. This approach can also be extended to the case when there are more than 2 constraints.

Example 1.5: Figure 1.13 shows a concrete example of the approximation discussed above for the unconstrained minimax reconstruction of (1.94). This reconstruction is compared with the LS solution $\hat{x}_{\mathrm{LS}} = \min_{\|Lx\| \leq \rho} \|S^* x - c\|^2$. In this experiment, the sampling filter is a rectangular window whose support is equal to the sampling interval, the noise is white and Gaussian with variance σ^2, and the signal $x = (x(1), \ldots, x(100))^T$ is given by $x(n) = \mathrm{sinc}(0.1(n - 33.3)) + \exp\{-0.005(n - 50.5)^2\}$. The regularization operator L was taken to be a (discrete approximation of) first order derivative. The parameter α was chosen as $3\sqrt{K}\sigma$, where K is the number of samples. Both the minimax and LS reconstructions were produced with the same ρ. It can be seen that $\hat{x}_{\mathrm{LS}}$ tends to oscillate more than $\hat{x}_{\mathrm{MX}}$. Consequently, its reconstruction error is larger than that of $\hat{x}_{\mathrm{MX}}$ by 30%.

To conclude, we have seen that when the samples are perturbed by noise, we can obtain an approximation of the minimax recovery by numerically solving an SDP. In some cases, this strategy leads to the exact minimax solution. The minimax method often yields improved reconstructions over LS. This is especially true in the constrained setting, where the minimax recovery is the orthogonal projection of the unconstrained solution onto the reconstruction space, whereas the LS recovery may deviate substantially from the unconstrained method.

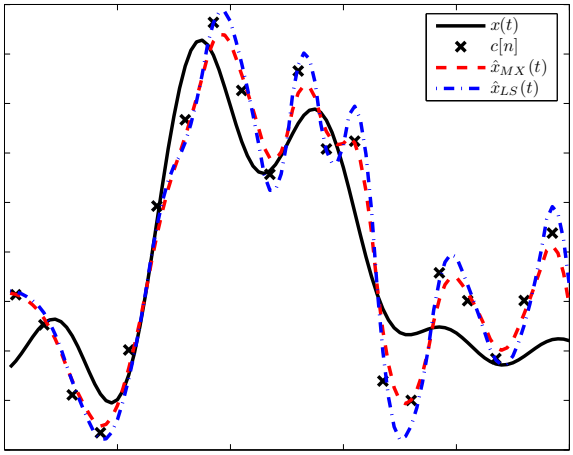

Figure 1.13: Comparison of minimax and LS reconstruction. The error norm of $\hat{x}_{\mathrm{LS}}$ is 30% higher than that of $\hat{x}_{\mathrm{MX}}$.

1.9 Conclusion

In this chapter we revisited the fundamental problem of reconstructing signals from their samples. We considered several models for each of the essential ingredients of the sampling problem: the sampling mechanism (general pre-filter, noise), the reconstruction kernel (pre-specified, unrestricted), and the signal prior (subspace, smoothness). Our approach was to define an optimization problem that takes into account both the fit of the reconstructed signal to the given samples and the prior knowledge we have about the signal. Each of the settings studied in this chapter was treated using two optimization strategies: LS and minimax. We showed that when the samples are noise-free, both strategies coincide if the reconstruction mechanism is unrestricted. In this case, perfect recovery is often possible under a subspace prior. In contrast, when the recovery is constrained, the minimax strategy leads to solutions that are closer to the original signal. The last part of this chapter was devoted to the challenging task of treating smoothness priors via the minimax strategy in the case in which the samples are noisy. Since closed-form solutions are unavailable in this setting, we restricted our attention to signals lying in $\mathbb{R}^n$ and $\mathbb{C}^n$ and showed how the resulting problems can be solved numerically using standard optimization packages. This was made possible by relying on recent results in optimization theory, regarding the tightness of SDP relaxations in quadratic problems.

Acknowledgments

We thank Dr. Ewa Matusiak for fruitful discussions.

This work was supported in part by the Israel Science Foundation under Grant no. 1081/07 and by the European Commission in the framework of the FP7 Network of Excellence in Wireless COMmunications NEWCOM++ (contract no. 216715).

References

[1] C. E. Shannon, "Communications in the presence of noise," *Proc. IRE*, vol. 37, pp. 10–21, Jan 1949.

[2] E. T. Whitaker, "On the functions which are represented by the expansion of interpolating theory," in *Proc. Roy. Soc. Edinburgh*, vol. 35, 1915, pp. 181–194.

[3] V. A. Kotelnikov, "On the transmission capacity of ether," in *Proc. First All-union Conference on Questions of Communications*, 1933.

[4] H. Nyquist, "Certain topics in telegraph transmission theory," *AIEE Trans.*, vol. 47, pp. 617–644, Jan. 1928.

[5] 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.

[6] A. Aldroubi and M. Unser, "Sampling procedures in function spaces and asymptotic equivalence with Shannon's sampling theory," *Numer. Funct. Anal. Optimiz.*, vol. 15, pp. 1–21, Feb. 1994.

[7] A. Aldroubi and K. Gröchenig, "Non-uniform sampling and reconstruction in shift-invariant spaces," *Siam Review*, vol. 43, pp. 585–620, 2001.

[8] Y. C. Eldar, "Sampling and reconstruction in arbitrary spaces and oblique dual frame vectors," *J. Fourier Analys. Appl.*, vol. 1, no. 9, pp. 77–96, Jan. 2003.

[9] ——, "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.

[10] I. J. Schoenberg, "Spline functions and the problem of graduation," *Proceedings of the National Academy of Sciences*, vol. 52, no. 4, pp. 947–950, 1964.

[11] S. Ramani, D. Van De Ville, T. Blu, and M. Unser, "Nonideal Sampling and Regularization Theory," *IEEE Trans. Signal Processing*, vol. 56, no. 3, pp. 1055–1070, 2008.

[12] M. Unser and T. Blu, "Generalized smoothing splines and the optimal discretization of the Wiener filter," *IEEE Trans. Signal Processing*, vol. 53, no. 6, pp. 2146–2159, 2005.

[13] T. Michaeli and Y. C. Eldar, "High Rate Interpolation of Random Signals from Nonideal Samples," *IEEE Trans. Signal Processing*, vol. 57, no. 3, pp.

977–992, March 2009.

[14] 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.

[15] Y. C. Eldar and T. Michaeli, "Beyond bandlimited sampling: nonlinearities, smoothness and sparsity," to appear in *IEEE Signal Processing Magazine*.

[16] S. Kayalar and H. L. Weinert, "Oblique projections: Formulas, algorithms, and error bounds," *Math. Contr. Signals Syst.*, vol. 2, no. 1, pp. 33–45, 1989.

[17] O. Christensen, *Frames and Bases. An Introductory Course*. Boston, MA: Birkhäuser, 2008.

[18] I. Daubechies, "The wavelet transform, time-frequency localization and signal analysis," *IEEE Trans. Inform. Theory*, vol. 36, pp. 961–1005, Sep. 1990.

[19] ——, *Ten Lectures on Wavelets*. SIAM, Philadelphia, 1992.

[20] R. M. Young, *An Introduction to Nonharmonic Fourier Series*. New York: Academic Press, 1980.

[21] M. Unser, "Sampling—50 years after Shannon," *IEEE Proc.*, vol. 88, pp. 569–587, Apr. 2000.

[22] M. Mishali and Y. C. Eldar, "Blind multi-band signal reconstruction: Compressed sensing for analog signals," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 993–1009, March 2009.

[23] Y. C. Eldar, "Compressed sensing of analog signals in shift invariant spaces," to appear in *IEEE Trans. on Signal Process.*

[24] D. L. Donoho, "Compressed sensing," *IEEE Trans. on Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr 2006.

[25] 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.

[26] M. Mishali and Y. C. Eldar, "Reduce and boost: Recovering arbitrary sets of jointly sparse vectors," *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 4692–4702, Oct. 2008.

[27] Y. C. Eldar, "Uncertainty relations for analog signals," submitted to *IEEE Trans. Inform. Theory*.

[28] Y. M. Lu and M. N. Do, "A theory for sampling signals from a union of subspaces," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2334–2345, 2008.

[29] Y. C. Eldar and M. Mishali, "Robust recovery of signals from a union of subspaces," submitted to *IEEE Trans. Inform. Theory*.

[30] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D*, vol. 60, no. 1-4, pp. 259–268, 1992.

[31] 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.

[32] 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.

[33] A. Aldroubi, "Oblique projections in atomic spaces," *Proc. Amer. Math. Soc.*, vol. 124, no. 7, pp. 2051–2060, 1996.

[34] T. G. Dvorkind, Y. C. Eldar, and E. Matusiak, "Nonlinear and non-ideal sampling: Theory and methods," *IEEE Trans. Signal Process.*, vol. 56, no. 12, pp. 5874–5890, 2008.

[35] R. G. Keys, "Cubic convolution interpolation for digital image processing," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 6, pp. 1153–1160, 1981.

[36] 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.

[37] I. J. Schoenberg, *Cardinal Spline Interpolation*. Philadelphia, PA: SIAM, 1973.

[38] 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.

[39] ——, "B-Spline signal processing: Part II - Efficient design and applications," *IEEE Trans. Signal Processing*, vol. 41, no. 2, pp. 834–848, Feb. 1993.

[40] Y. C. Eldar, "Rethinking biased estimation: Improving maximum likelihood and the Cramer–Rao bound," *Foundations and Trends in Signal Processing*, vol. 1, no. 4, pp. 305–449, 2007.

[41] Y. C. Eldar, A. Ben-Tal, and A. Nemirovski, "Robust mean-squared error estimation in the presence of model uncertainties," *IEEE Trans. Signal Processing*, vol. 53, no. 1, pp. 168–181, Jan. 2005.

[42] 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.

[43] M. Sion, "On general minimax theorems," *Pac. J. Math.*, vol. 8, pp. 171–176, 1958.

[44] A. Beck and Y. C. Eldar, "Strong duality in nonconvex quadratic optimization with two quadratic constraints," *Siam J. Optimization*, vol. 17, no. 3, pp. 844–860, 2006.

[45] Y. C. Eldar and O. Christansen, "Characterization of oblique dual frame pairs," *J. Applied Signal Processing*, pp. 1–11, 2006, article ID 92674.

[46] O. Christansen and Y. C. Eldar, "Oblique dual frames and shift-invariant spaces," *Appl. Comp. Harm. Anal.*, vol. 17, no. 1, pp. 48–68, 2004.

[47] 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.

[48] Y. C. Eldar and N. Merhav, "A competitive minimax approach to robust estimation of random parameters," *IEEE Trans. Signal Processing*, vol. 52, pp. 1931–1946, July 2004.

[49] E. Lieb and M. Loss, *Analysis*. American Mathematical Society, 2001.
[50] A. Ben-Tal and M. Teboulle, "Hidden convexity in some nonconvex quadratically constrained quadratic programming," *Mathematical Programming*, vol. 72, no. 1, pp. 51–63, 1996.
[51] H. G. Feichtinger and T. Werther, "Robustness of minimal norm interpolation in sobolev algebras," in *Sampling, Wavelets and Tomography*, A. I. Zayed and J. J. Benedetto, Eds. Boston, MA: Birkhäuser, 2004.
[52] J. M. Martínez, "Local minimizers of quadratic functions on Euclidean balls and spheres," *SIAM J. Optim.*, vol. 4, no. 1, pp. 159–176, 1994.
[53] J. J. Moré and D. C. Sorensen, "Computing a trust region step," *SIAM J. Sci. Statist. Comput.*, vol. 4, no. 3, pp. 553–572, 1983.
[54] D. C. Sorensen, "Newton's method with a model trust region modification," *SIAM J. Numer. Anal.*, vol. 19, no. 2, pp. 409–426, 1982.
[55] R. J. Stern and H. Wolkowicz, "Indefinite trust region subproblems and nonsymmetric eigenvalue perturbations," *SIAM J. Optim.*, vol. 5, no. 2, pp. 286–313, 1995.
[56] L. Vandenberghe and S. Boyd, "Semidefinite programming," *SIAM Rev.*, vol. 38, no. 1, pp. 40–95, Mar. 1996.
[57] M. R. Celis, J. E. Dennis, and R. A. Tapia, "A trust region strategy for nonlinear equality constrained optimization," in *Numerical optimization, 1984 (Boulder, Colo., 1984)*. Philadelphia, PA: SIAM, 1985, pp. 71–82.
[58] A. R. Conn, N. I. M. Gold, and P. L. Toint, *Trust-Region Methods*, ser. MPS/SIAM Series on Optimization. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM), 2000.
[59] M. J. D. Powell and Y. Yuan, "A trust region algorithm for equality constrained optimization," *Math. Programming*, vol. 49, no. 2, (Ser. A), pp. 189–211, 1990/91.
[60] Y. Yuan, "On a subproblem of trust region algorithms for constrained optimization," *Mathematical Programming*, vol. 47, pp. 53–63, 1990.
[61] Y. M. Zhu, "Generalized sampling theorem," *IEEE Tran. Circ. Sys. II*, vol. 39, pp. 587–588, Aug 1992.
[62] Y. Ye and S. Zhang, "New results on quadratic minimization," *SIAM J. Optim.*, vol. 14, no. 1, pp. 245–267, 2003.
[63] A. L. Fradkov and V. A. Yakubovich, "The S-procedure and the duality relation in convex quadratic programming problems," *Vestnik Leningrad. Univ.*
[64] N. Z. Shor, "Quadratic optimization problems," *Izv. Akad. Nauk SSSR Tekhn. Kibernet.*, no. 1, pp. 128–139, 222, 1987.
[65] A. Beck and Y. C. Eldar, "Regularization in regression with bounded noise: A Chebyshev center approach," *SIAM J. Matrix Anal. Appl.*, vol. 29, no. 2, pp. 606–625, 2007.
[66] Y. C. Eldar, A. Beck, and M. Teboulle, "A minimax Chebyshev estimator for bounded error estimation," *IEEE Trans. Signal Processing*, vol. 56, no. 4,

pp. 1388–1397, Apr. 2008.

[67] S. Boyd, L. E. Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*. Philadelphia, PA: SIAM, 1994.

[68] J. F. Sturm, "Using sedumi 1.02, a matlab toolbox for optimization over symmetric cones," *Optimization Methods and Software*, vol. 11-12, pp. 625–653, 1999.

[69] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming (web page and software)," March 2008, `http://stanford.edu/~boyd/cvx`.

Index