Trading Accuracy for Numerical Stability:
Orthogonalization, Biorthogonalization and Regularization

Tarek A. Lahlou and Alan V. Oppenheim
Digital Signal Processing Group
Massachusetts Institute of Technology

Abstract—This paper presents two novel regularization methods motivated in part by the geometric significance of biorthogonal bases in signal processing applications. These methods, in particular, draw upon the structural relevance of orthogonality and biorthogonality principles and are presented from the perspectives of signal processing, convex programming, continuation methods and nonlinear projection operators. Each method is specifically endowed with either a homotopy or tuning parameter to facilitate tradeoff analysis between accuracy and numerical stability. An example involving a basis comprised of real exponential signals illustrates the utility of the proposed methods on an ill-conditioned inverse problem and the results are compared to standard regularization techniques from the signal processing literature.

Index Terms—orthogonality, biorthogonality, regularization, convex optimization, homotopy maps

I. INTRODUCTION

Signal processing algorithms, notably those in machine learning and big data applications, make ubiquitous use of dense and sparse linear algebra subroutines [1], [2]. Numerical errors often arise while implementing such subroutines and are typically explained using perturbation analysis tools. For example, finite-precision effects such as coefficient quantization are quantifiable by invoking sensitivity theorems after organizing the computation into a linear or nonlinear signal-flow graph, as discussed in, e.g., [3]. Also, the magnification and compounding of propagated errors are commonly understood via the condition number of a functional representation of the processing system. These unavoidable sources of error contribute significantly to the total accumulated error in an obtained solution, especially as problem sizes continue to scale. Algorithms specifically designed to use orthogonality principles alleviate these issues and, among other reasons, receive extensive use in practice. Moreover, orthogonality in other contexts such as quantum measurement and asynchronous distributed optimization continue to provide inspiration for new signal processing algorithms, architectures, and applications [4], [5].

Signal processing techniques exploiting natural and efficient representations for signal models with inherent geometric structure often use non-orthogonal bases and frames that may generally suffer from conditioning issues, e.g. finite rate of innovation [6], compressed sensing [7], and transient signal analysis [8]-[11]. To address these issues, we restrict our attention in this paper to two classical approaches: orthogonalization and regularization. Traditional orthogonalization routines typically destroy the topological relevance of an original basis to the problem at hand while simultaneously weakening the informational significance of the solution obtained. The orthogonal Procrustes problem has been used for subspace whitening and machine learning in the matrix setting, i.e. generating the unitary matrix that optimally maps one matrix to another in the Frobenius norm sense[12], [13]. In this paper, we specifically translate and explicate these results using bases and further extend them as

they pertain to biorthogonalization, nonlinear projections, and novel regularization methods. The proposed regularization methods permit natural characterizations of accuracy and stability enabling standard tradeoff analysis to identify an acceptable balance. Regularization methods in signal processing have historically focused on directly obtaining a solution biased toward prespecified signal properties. In this paper, however, we emphasize the effect the proposed methods have on the linear problems Euclidean structure itself.

This paper is outlined as follows: In Section II we collect together geometric facts surrounding orthogonality and biorthogonality that will be referenced throughout. The orthogonalization and biorthogonalization of a linear system via a nonlinear projection operator is the central focus of Section III. Two regularization methods for balancing numerical robustness in the sense of relative conditioning and accuracy in the sense of geometric structure are developed in Section IV. The proposed methods are then applied to the regularization of an ill-conditioned basis composed of real exponential signals in Section V as a numerical example. Throughout the presentation, we draw no distinction between various linear problems and the mathematical relationships connecting them, e.g. matched filtering, solving linear systems, changing bases, etc. In this way, the insights and methods presented may be readily applied to any linear problem equally well.

A. Notational conventions

Vectors are denoted using underscores with subscripts indexing vectors as opposed to entries, i.e. $\underline{a}_1$ and $\underline{a}_2$ represent two distinct vectors. Boldface letters denote a *system* or collection of N vectors, typically constituting a basis for $\mathbb{R}^N$, i.e. $\mathbf{a} = \{\underline{a}_k\}_{k=1}^N$. A biorthogonal set is denoted using $\sim$, i.e. $\widetilde{\mathbf{a}}$ is biorthogonal to $\mathbf{a}$. The p-norm of a vector $\underline{a}$ for $p \geq 1$ is denoted $\|\underline{a}\|_p$. A capital letter denotes the matrix form of a set of vectors, i.e. A has columns $\mathbf{a}$. The Kronecker delta is denoted $\delta_{k,j}$. For clarity we assume all systems are real-valued bases; rigorous extensions to complex-valued bases and frames follow analogously, e.g. by using appropriate conjugation and restricting arguments to the range of a system.

II. THE TOPOLOGY OF THE STIEFEL MANIFOLD AND $O(N)$

Let $\mathbf{a}$ and $\mathbf{b}$ denote two general systems and define ϵ_{BO} as the nonnegative measure of their biorthogonality given by

$$\epsilon_{BO}(\mathbf{a}, \mathbf{b}) \triangleq \sum_k \sum_j \left(\langle \underline{a}_k, \underline{b}_j \rangle - \delta_{k,j} \right)^2 \qquad (1)$$

which achieves a unique global minimum of zero provided that $\mathbf{b} = \widetilde{\mathbf{a}}$. It is then straightforward to conclude that

$$\epsilon_{BO}(\mathbf{a}, \mathbf{a}) = 0 \iff \mathbf{a} \text{ is an orthonormal system.} \qquad (2)$$

The Stiefel manifold $V_{k,N}$ is the orthogonal groups principal homogenous space $(k = N)$ and is written using (1) as [14]

$$V_{N,N} = \left\{ \mathbf{h} \in \mathbb{R}^N : \epsilon_{BO}(\mathbf{h}, \mathbf{h}) = 0 \right\}. \qquad (3)$$

Represented as a quadratic form, (1) is convex quadratic with positive signature, i.e. all eigenvalues are positive valued, thus any

optimization problem that uses (1) as a cost function or (3) as a feasible set is not precluded from being a convex problem.

The orthogonal group of dimension N, denoted by $O(N)$ and described in set-builder notation using the matrix form of **h** as

$$O(N) = \left\{ H \in \mathbb{R}^{N \times N} : H^T H = H H^T = I_N \right\} \qquad (4)$$

where I_N is the identity operator on $\mathbb{R}^N$, consists of two connected components $\mathcal{M}_1$ and $\mathcal{M}_{-1}$. Furthermore, $\mathcal{M}_1$ and $\mathcal{M}_{-1}$ each form a disjoint smooth manifold in $\mathbb{R}^{N \times N}$ characterized by

$$\mathcal{M}_i = \{ H \in O(N) : \det(H) = i \}, \quad i = \pm 1 \qquad (5)$$

where $\det(\cdot)$ denotes the determinant operator. Note that $\mathcal{M}_1$ is precisely the special orthogonal group $SO(N)$. From (5) and the definition of a connected component it immediately follows that

$$H \in \mathcal{M}_i \Leftrightarrow \widetilde{H} \in \mathcal{M}_i, \quad i = \pm 1. \qquad (6)$$

We conclude this section by noting that appropriately selecting a base point for $V_{N,N}$ establishes a one-to-one correspondence with $O(N)$ and, therefore, justifies interchanging between the two notational descriptions of an orthogonal system, i.e. **h** and H.

III. Orthogonalization and Biorthogonalization

Let **a** and **b** continue to denote two general systems and define ϵ_{LS} as the nonnegative Euclidean measure of their distance given by

$$\epsilon_{LS}(\mathbf{a}, \mathbf{b}) \triangleq \sum_k \| \underline{a}_k - \underline{b}_k \|_2^2 \qquad (7)$$

$$= \sum_k \langle \underline{a}_k - \underline{b}_k, \underline{a}_k - \underline{b}_k \rangle \qquad (8)$$

which achieves a unique global minimum of zero provided that $\mathbf{b} = \mathbf{a}$. In the sequel we shall denote by **g** a general prespecified system.

A. Optimal orthogonalization problems

Consider the optimization problem that identifies the nearest orthogonal system to **g** in the sense of (7), or equivalently the least-squares orthogonalization of **g**, given by

$$\widehat{\mathbf{h}}_1 = \arg\min_{\mathbf{h}} \epsilon_{LS}(\mathbf{g}, \mathbf{h}) \ \text{ s.t. } \epsilon_{BO}(\mathbf{h}, \mathbf{h}) = 0. \qquad (P1)$$

Similarly, consider the optimization problem that identifies the nearest orthogonal system to the biorthogonal system $\widetilde{\mathbf{g}}$ in the sense of (7) given by

$$\widehat{\mathbf{h}}_2 = \arg\min_{\mathbf{h}} \epsilon_{LS}(\widetilde{\mathbf{g}}, \mathbf{h}) \ \text{ s.t. } \epsilon_{BO}(\mathbf{h}, \mathbf{h}) = 0. \qquad (P2)$$

We collectively refer to (P1) and (P2) as optimal orthogonalization problems and further interpret their solutions as orthogonal systems that maximally preserve the geometric structure inherent to **g** and $\widetilde{\mathbf{g}}$, respectively.

B. Optimal biorthogonalization problems

Consider the optimization problem that identifies the orthogonal system which is maximally dual to **g** in the sense of (1) given by

$$\widehat{\mathbf{h}}_3 = \arg\min_{\mathbf{h}} \epsilon_{BO}(\mathbf{g}, \mathbf{h}) \ \text{ s.t. } \epsilon_{BO}(\mathbf{h}, \mathbf{h}) = 0. \qquad (P3)$$

Finally, consider the optimization problem that identifies the orthogonal system which is maximally dual to the biorthogonal system $\widetilde{\mathbf{g}}$ in the sense of (1) given by

$$\widehat{\mathbf{h}}_4 = \arg\min_{\mathbf{h}} \epsilon_{BO}(\widetilde{\mathbf{g}}, \mathbf{h}) \ \text{ s.t. } \epsilon_{BO}(\mathbf{h}, \mathbf{h}) = 0. \qquad (P4)$$

We collectively refer to (P3) and (P4) as optimal biorthogonalization problems and further interpret their solutions as the orthogonal systems that, switching to the matrix form notation introduced in Subsection I-A, maximally invert G and $\widetilde{G}$, respectively.
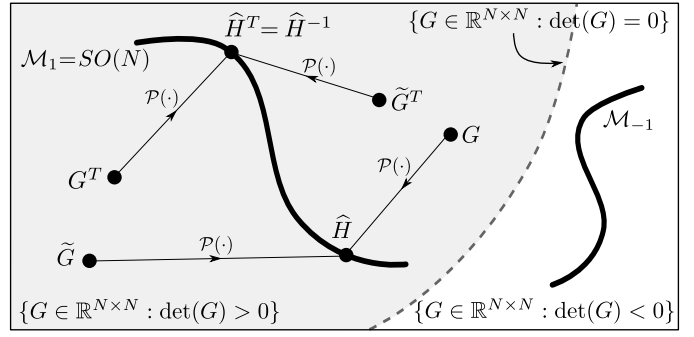

Fig. 1. An illustration of the relationships between the matrix forms of **g** and $\widetilde{\mathbf{g}}$ and their adjoints and their projections onto the orthogonal group via (13) where the projection operator is denoted using $\mathcal{P}(\cdot)$. The system **g** is chosen such that $\det(G) > 0$ and therefore the projections are onto $\mathcal{M}_1$.

C. Equivalence of ϵ_{LS} and ϵ_{BO} over $V_{N,N}$

The measures (1) and (7), respectively interpreted as optimal biorthogonalization and orthgonalization cost functions, produce generally unequal cost surfaces over the space of all N-dimensional bases. However, for fixed **g** and orthogonal **h**, it follows that

$$\epsilon_{BO}(\mathbf{g}, \mathbf{h}) = \sum_{k,j} \left(\langle \underline{g}_k, \underline{h}_j \rangle \left[\langle \underline{g}_k, \underline{h}_j \rangle - 2\delta_{k,j} \right] + \delta_{k,j}^2 \right) \qquad (9)$$

$$= \sum_k \left(\langle \underline{g}_k, \underline{g}_k \rangle - 2\langle \underline{g}_k, \underline{h}_k \rangle + \langle \underline{h}_k, \underline{h}_k \rangle \right) \qquad (10)$$

$$= \sum_k \langle \underline{g}_k - \underline{h}_k, \underline{g}_k - \underline{h}_k \rangle \qquad (11)$$

$$= \epsilon_{LS}(\mathbf{g}, \mathbf{h}) \qquad (12)$$

where the restriction of **h** to $V_{N,N}$ is used to obtain (10) from (9), in particular since $\langle \underline{a}, \underline{b} \rangle = \sum_j \langle \underline{a}, \underline{h}_j \rangle \langle \underline{h}_j, \underline{b} \rangle$ and $\delta_{k,j} = \langle \underline{h}_k, \underline{h}_j \rangle$.

D. Analytic solution and projection interpretation of (P1)-(P4)

In Subsection III-C, we established that the solution to (P1) solves (P3) and likewise that the solution to (P2) solves (P4). In this subsection, we first state and interpret the analytic expression that minimizes (P1) followed by discussing one way (P1) and (P2) may be argued to have the same minimizer. Indeed, the matrix form of the orthogonal system $\widehat{\mathbf{h}}_1$ which solves (P1) is given by

$$\widehat{H}_1 = U V^T \qquad (13)$$

where $G = U \Sigma V^T$ is the Singular Value Decomposition (SVD) of the matrix form of **g**. Observe that verifying the orthogonality of (13) folllows immediately from the closure property of $O(N)$; a sketch of the full argument that (13) minimizes (P1) is deferred to the Appendix. Interpreting (13) as a nonlinear operator acting upon G, it is straightforward to verify idempotency, hence (13) is the projection of G onto $O(N)$. Figure 1 illustrates this, i.e. that the action of (13) on G, $\widetilde{G}$ and their transposes is a projection onto the orthogonal group. From the perspective of bases, generating (13) consists of N radial projections along the semiaxes of the ellipsoid described by the SVD of G such that the associated singular values become unity.

An equivalence between (P1) and (P2) is established by using the matrix and basis representation equivalence discussed earlier where the matrix form of $\widetilde{\mathbf{g}}$, i.e. the linear functionals biorthogonal to **g** in the dual linear space, corresponds to the inverse of the adjoint of G. The desired result is obtained by straightforward manipulations that further simplify (P2), in particular by exploiting SVD properties. For arbitrary systems **g**, we comment that it is possible to identify *a priori* which manifold (13) belongs to, i.e. combining the manifold

characterizations in (5) and using the connected component property in (6) we conclude that

$$\widehat{H} \in \mathcal{M}_i \quad \Longleftrightarrow \quad \mathrm{sgn}\left(\det(G)\right) = i, \quad i = \pm 1. \qquad (14)$$

E. Incremental projection methods for nearly orthogonal **g**

Although many numerical algorithms exist to compute (13), we call special attention to the case of G being "nearly" orthogonal. Let $R = G^T G - I_N$ be a Grammian residual. Re-writing (13) as

$$\widehat{H} = G(I_N + R)^{-\frac{1}{2}}, \qquad (15)$$

we proceed assuming the spectral norm of R is less than 1 which may be efficiently certified using, e.g. Gershgorin's circle theorem [15]. After some straightforward manipulations we obtain the power series

$$\widehat{H} = G \left[\sum_{n=0}^{\infty} \binom{-\frac{1}{2}}{n} R^n \right] \qquad (16)$$

and thus the solution to (P1)-(P4) may be generated to any desired level of accuracy without explicitly computing an SVD by truncating (16) to an appropriate number of terms.

IV. Trading accuracy for numerical stability

Many applications of signal processing as well as engineering more broadly require considerably less accuracy than what is available on modern computational platforms. By pairing systems in $V_{N,N}$ with backward stable algorithms, computation can be made to not magnify propagated errors. Motivated by these observations, we next propose two regularization methods that facilitate accuracy and stability tradeoffs readily explained through homotopy maps and convex analysis, respectively. We comment upfront that while our treatment focuses on interpretation, the utility of these methods is in regularizing ill-conditioned problems while preserving some of the original systems Euclidean structure as is illustrated in Section V.

Regularization methods are often described initially using the language of optimization theory; the methods surveyed in Sub-section V-B serve as examples. Some regularization techniques, primarily those related to least squares formulations, admit solutions in closed-form. Setting these special cases aside, the remaining techniques conventionally solve directly for a solution to the problem at hand without first generating a regularized system. The methods proposed in this section specifically make this intermediary step, i.e. they explicitly generate a regularized linear system from which a solution is then generated, potentially making further use of any regularization techniques belonging to the class just mentioned in which the regularized system takes the place of the original system.

A. Homotopic continuation formulation

Motivated by the extensive use of floating-point arithmetic in big data computation, we measure the stability of a system **g** through its relative condition number, i.e. the ratio of extremal singular values of G. Loosely speaking, when this measure is on the order 10^k then only $d - k$ digits of the computed solution are reliable where d is the maximum number available on the chosen computational platform [16]. This loss in accuracy is inherent to the system itself and says nothing about further degradations that accumulate due to a particular algorithms stability nor to errors that arise from round-off noise and coefficient quantization.

We first introduce a continuation scheme to generate a parameterized family of regularized systems along the trajectory corresponding to the smooth deformation of **g** into $\widehat{\mathbf{h}}_1$. The utility of this scheme for solving a linear system of equations $G\underline{x} = \underline{y}$, with known $\underline{y}$,

is to sequentially or approximately solve for $\underline{x}$ while substituting different systems along this continuum for G. The parameterization of the system mitigates numerical conditioning issues at the expense of accuracy, i.e. the continuation terminates when a system associated with an acceptable solution is identified. The function describing these deformations is referred to as a *homotopy map*. In this paper, we use the specific homotopy map

$$\widehat{\mathbf{h}}_5(\rho, \mathbf{g}) \triangleq \left\{ \underline{h}_k : \underline{h}_k = (1-\rho)\underline{g}_k + \rho \underline{z}_k \text{ where } \mathbf{z} \text{ solves (P1)} \right\} \quad (17)$$

where the parameter $\rho \in [0,1]$ balances the previously addressed tradeoff. The family of linear systems achieved using the homotopy map $\widehat{\mathbf{h}}_5$ is $\widehat{\mathbf{h}}_5(\rho, \mathbf{g}) \leftrightarrow \widehat{H}_5$ for $\rho \in [0,1]$ and the standard implementation strategy is to track the solution to the problem using $\widehat{H}_5$ instead of G starting from $(\rho, \mathbf{g}) = (0, \mathbf{g})$ as ρ progresses from 0 to $\rho^\star$ where $\rho^\star$ generates the system associated with an acceptable solution. In Section V, solutions to a similar problem are obtained by discretizing $[0,1]$ and evaluating a performance metric for all values of ρ in this interval. Standard line search methods help to efficiently identify the optimal homotopy value $\rho^\star$.

Since the condition number of G is equal to the condition number of $\widetilde{G}$, and with the equivalence between the minimizer of (P1) and (P2), we conclude the discussion in this subsection by noting that replacing **g** with $\widetilde{\mathbf{g}}$ results in the same presentation with the geometric structure of the biorthogonal system $\widetilde{\mathbf{g}}$ being preserved instead.

B. Unconstrained optimization formulation

Prompted by the exposition on characterizing numerical accuracy and stability in the previous subsection, we next formulate an optimization problem whose cost function is the sum of the metrics (7) and (1), i.e. each summand is respectively minimized by $\widehat{\mathbf{h}}_5(0, \mathbf{g})$ and $\widehat{\mathbf{h}}_5(1, \mathbf{g})$. We formally write this problem as an unconstrained quartic convex optimization problem of the form

$$\widehat{\mathbf{h}}_6(\rho, \mathbf{g}) \triangleq \arg\min_{\mathbf{h}} \epsilon_{LS}(\mathbf{g}, \mathbf{h}) + \rho \epsilon_{BO}(\mathbf{h}, \mathbf{h}) \qquad (18)$$

where ρ again determines the previously addressed tradeoff. Indeed, consistent with the discussion surrounding the positive definite nature of (1) represented as a quadratic form and the fact that the sum of unconstrained convex functions results in a convex function, it is straightforward to verify that any solution to (18) that is locally optimal is also globally so. Although the cost function of (18) is purposefully constructed using the endpoints of the continuum of regularized systems achievable using (17), it is readily possible to generate solutions to (18) that do not lie on this continuum nor are obtainable using (17) for any value of the homotopy parameter ρ.

In response to the global optimality guarantees established above and for the sake of completeness, we conclude this section by including gradient primitives that may be used by any number of gradient-based nonlinear programming algorithms in solving (18):

$$\nabla \epsilon_{LS}(\mathbf{g}, \mathbf{h}) \mid_{\underline{h}_k} = 2 \sum_j (\langle \underline{g}_j, \underline{h}_k \rangle - \delta_{j,k}) \underline{g}_j \qquad (19)$$

$$\nabla \epsilon_{BO}(\mathbf{h}, \mathbf{h}) \mid_{\underline{h}_k} = 4 \left(\langle \underline{h}_k, \underline{h}_k \rangle - 1 \right) \underline{h}_k + 2 \sum_{j \neq k} \langle \underline{h}_j, \underline{h}_k \rangle \underline{h}_j. \quad (20)$$

Note that an analytic expression for the system $\widehat{\mathbf{h}}_6$ which minimizes (18) is currently unknown due to the cubic nature of the first order optimality conditions, i.e. the nonlinear system of equations defined by setting the appropriately weighted sum of the two terms above equal to zero for each value of k.

V. NUMERICAL EXAMPLE

Solving linear equations using regularization, a common approach to ill-posed problems, modifies the problem so that (i) the new problem is biased toward expected solutions and (ii) the previously mentioned numerical issues are reduced. In signal processing, regularization appears in many forms: Tikhonov regularization and Wiener filtering, total variation denoising in image processing, and basis pursuit denoising in compressive sensing. In this section, we use the methods developed in Section IV to solve a numerical example.

A. Problem formulation

Let E denote the matrix form of a real exponential basis $\mathbf{e}$, i.e. E is Vandermonde and, for $0 < \sigma_1 < \cdots < \sigma_N < 1$, generated by

$$E_{i,j} = \sigma_j^{i-1}, \quad 1 \le i, j \le N. \tag{21}$$

The relative condition number of E, and consequently $\widetilde{E}$ too, grows exponentially in N [17]. This fact justifies using regularization to solve a linear problem of the form

$$E\underline{x} = \underline{y} \tag{22}$$

for $\underline{x}$ where $\underline{y}$ is analytically generated according to a synthetic solution $\underline{\bar{x}}$ to help limit observed numerical effects to the regularization effect each method has on the matrix E in (22).

B. Numerical methods for comparison

This subsection briefly reviews three regularization methods from the signal processing literature, two of which originate in the compressive sensing setting, to compare against (17) and (18) in solving (22) [7]. The first method, Tikhonov regularization, is formulated as

$$\underline{x}^\star = \arg\min_{\underline{x}} \|E\underline{x} - \underline{y}\|_2^2 + \|T\underline{x}\|_2^2. \tag{23}$$

where we proceed to take the Tikhonov matrix $T = \rho I_N$. The objective in (23) is convex quadratic, thus the solution has an analytic expression, namely $\underline{x}^\star = (E^T E + \rho^2 I)^{-1} E^T \underline{y}$. Next, the basis pursuit denoising problem is formulated as

$$\underline{x}^\star = \arg\min_{\underline{x}} \|E\underline{x} - \underline{y}\|_2 + \rho\|\underline{x}\|_1 \tag{24}$$

where ρ balances the absolute size of the solution with the desired agreement of (22). Finally, the Dantzig selector is given by

$$\underline{x}^\star = \arg\min_{\underline{x}} \|E^T(E\underline{x} - \underline{y})\|_\infty + \rho\|\underline{x}\|_1 \tag{25}$$

and may be solved using standard linear programming techniques.

C. Numerical results

Our experimental setup is as follows: We generate E on each trial by selecting $N = 18$ values σ_k uniformly at random from the interval $[0.1, 0.9]$ and solve (22) using the various methods discussed. We select ρ to within 10^{-6} for each method per trial to minimize the residual $\|\underline{x}^\star - \underline{\bar{x}}\|_2$. The solution $x^\star$ to (17) and (18) is found by directly solving the regularized equations via Gaussian elimination. Table 1 lists the average residual and value of ρ taken over 10^7 trials. Solving (22) directly yields an average error of 8.64. The condition number of E for this experiment is lower bounded by 10^{16}. In summary, we remark that (17) outperforms the alternatives and (18) outperforms all but (23) for the Vandermonde linear systems tested. The larger error for (25) is due in part to the condition number of $E^T E$ being the square of the condition number of E.

Figure 2 illustrates an accuracy versus numerical stability tradeoff curve for the proposed regularization methods where accuracy is given by the residual $\|\underline{x}^\star - \underline{\bar{x}}\|_2$ and numerical stability is given by the
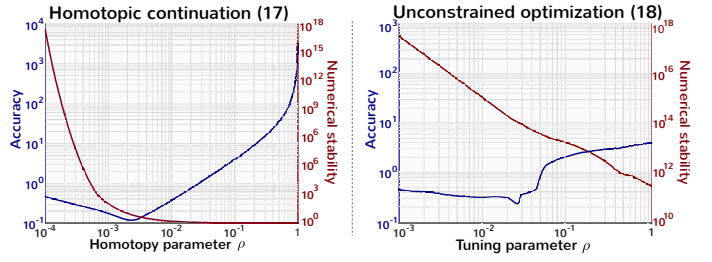

Fig. 2. Accuracy (residual error) versus numerical stability (relative condition number) tradeoff curves for (17) and (18).
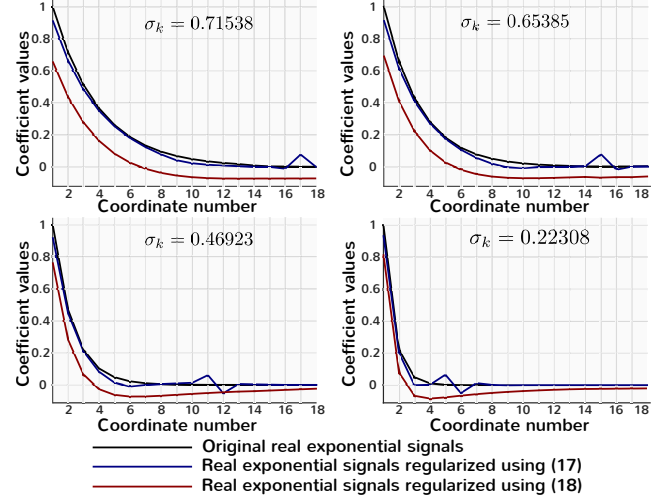

Fig. 3. An illustration of a subset of the real exponential basis elements $\mathbf{e}$ and the corresponding elements from the proposed regularization methods.

relative condition number. For both methods, it is clear that a non-zero value of ρ is optimal corresponding to a system whose conditioning is several orders of magnitude less than E's. Figure 3 depicts a subset of regularized signals for the optimal selection of ρ in Figure 2. As depicted and consistent with elements not shown, the regularized signals have similar structure to $\mathbf{e}$, i.e. they decay at similar rates to their exponential counterparts, also depicted for comparison. Signals produced by (18) retain a smooth nature with an additive offset and slight upward curvature and those produced by (17) contain small fluctuations that seem to alleviate conditioning issues.

TABLE 1. NUMERICAL COMPARISON OF REGULARIZATION METHODS

Solution reference	(17)	(18)	(23)	(24)	(25)
average $\|\underline{x}^\star - \underline{\bar{x}}\|_2$	0.096	0.153	0.139	0.147	13.46
average optimal ρ	0.002	0.028	0.025	0.015	0.044

APPENDIX

Sketch of the argument that (13) solves (P1):

$$\arg\min_{\mathbf{h}\in V_{N,N}} \epsilon_{LS}(\mathbf{g},\mathbf{h}) = \arg\min_{\mathbf{h}\in V_{N,N}} \sum_k \langle \underline{g}_k - \underline{h}_k, \underline{g}_k - \underline{h}_k \rangle$$

$$= \min_{\mathbf{h}\in V_{N,N}} \sum_k \left(\langle \underline{g}_k, \underline{g}_k \rangle + \langle \underline{h}_k, \underline{h}_k \rangle - 2\langle \underline{g}_k, \underline{h}_k \rangle \right)$$

$$= \arg\max_{\mathbf{h}\in V_{N,N}} \langle \underline{g}_1, \underline{h}_1 \rangle + \cdots + \langle \underline{g}_N, \underline{h}_N \rangle$$

Switching to matrix notation for convenience, utilizing properties of a trace, and writing G using its SVD $G = U\Sigma V^T$, we proceed as:

$$\arg\max_{\mathbf{h}\in V_{N,N}} \sum_k \langle \underline{g}_k, \underline{h}_k \rangle \rightarrow \arg\max_{H\in O(N)} \text{trace}\left(V\Sigma U^T H \right)$$

$$= U\left(\arg\max_{H\in O(N)} \text{trace}(\Sigma H) \right) V^T$$

$$= UV^T$$

REFERENCES

[1] K. Slavakis, G. B. Giannakis, and G. Mateos, "Modeling and optimization for big data analytics: (statistical) learning tools for our era of data deluge," *IEEE Signal Processing Magazine*, vol. 31, no. 5, pp. 18–31, Sept 2014.

[2] D. Yu and L. Deng, "Deep learning and its applications to signal and information processing," *IEEE Signal Processing Magazine*, vol. 28, no. 1, pp. 145–154, Jan 2011.

[3] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1975.

[4] Y. C. Eldar and A. V. Oppenheim, "Quantum signal processing," *IEEE Signal Processing Magazine*, vol. 19, no. 6, pp. 12–32, Nov 2002.

[5] T. A. Baran and T. A. Lahlou, "Conservative signal processing architectures for asynchronous, distributed optimization part I: General framework," in *IEEE Global Conference on Signal and Information Processing*, Dec 2014, pp. 35–39.

[6] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, Jun 2002.

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

[8] D. Potts and M Tasche, "Parameter estimation for nonincreasing exponential sums by pronylike methods," *Linear Algebra Applications*, vol. 439, pp. 1024–1039, 2013.

[9] T. A. Lahlou and A. V. Oppenheim, "Spectral representation of transient signals," in *IEEE Global Conference on Signal and Information Processing*, Dec 2014, pp. 660–663.

[10] V. Pereyra and G. Scherer, *Exponential Data Fitting and its Applications*, Bentham Science Publishers, Sharjah, 2010.

[11] L. Coluccio, A. Eisinberg, and G. Fedele, "A prony-like method for non-uniform sampling," *Signal Processing*, vol. 87, no. 10, pp. 2484 – 2490, 2007.

[12] S. Zafeiriou and N. Laskaris, "On the improvement of support vector techniques for clustering by means of whitening transform," *IEEE Signal Processing Letters*, vol. 15, pp. 198–201, 2008.

[13] Y. C. Eldar and A. V. Oppenheim, "MMSE whitening and subspace whitening," *IEEE Transactions on Information Theory*, vol. 49, no. 7, pp. 1846–1851, July 2003.

[14] A. Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, New York, 2002.

[15] R. S. Varga, *Gershgorin and His Circles*, ser. Springer Series in Computational Mathematics. Berlin, Germany: Springer-Verlag, 2004.

[16] W. Cheney and D. Kincaid, *Numerical Mathematics and Computing*, Thomson Brooks/Cole, 2007.

[17] B. Beckermann, "The condition number of real vandermonde, krylov and positive definite hankel matrices," *Numerische Mathematik*, vol. 85, pp. 553–577, 1997.