

Generalized SURE for Exponential Families: Applications to Regularization

Yonina C. Eldar, *Senior Member, IEEE*

Abstract—Stein’s unbiased risk estimate (SURE) was proposed by Stein for the independent, identically distributed (i.i.d.) Gaussian model in order to derive estimates that dominate least squares (LS). Recently, the SURE criterion has been employed in a variety of denoising problems for choosing regularization parameters that minimize an estimate of the mean-squared error (MSE). However, its use has been limited to the i.i.d. case which precludes many important applications. In this paper we begin by deriving a SURE counterpart for general, not necessarily i.i.d. distributions from the exponential family. This enables extending the SURE design technique to a much broader class of problems. Based on this generalization we suggest a new method for choosing regularization parameters in penalized LS estimators. We then demonstrate its superior performance over the conventional generalized cross validation and discrepancy approaches in the context of image deblurring and deconvolution. The SURE technique can also be used to design estimates without predefining their structure. However, allowing for too many free parameters impairs the estimate’s performance. To address this inherent tradeoff, we propose a regularized SURE objective, and demonstrate its use in the context of wavelet denoising.

Index Terms—Deconvolution, MSE estimation, regularization, Stein unbiased risk estimate.

I. INTRODUCTION

ESTIMATION in multivariate problems is a fundamental topic in statistical signal processing. One of the most common recovery strategies for deterministic unknown parameters is the well-known maximum-likelihood (ML) method. The ML estimator enjoys several appealing properties, including asymptotic efficiency under suitable regularity conditions. Nonetheless, its mean-squared error (MSE) can be improved upon in the non-asymptotic regime in many different settings.

In their seminal work, Stein and James showed that for the independent, identically distributed (i.i.d.) linear Gaussian model, it is possible to construct a nonlinear estimator with lower MSE than that of ML for all values of the unknowns [1], [2]. Various modifications of the James–Stein method have since been developed that are applicable to the non-i.i.d. Gaussian case as well [3]–[8].

Manuscript received November 01, 2007; revised August 09, 2008. First published October 31, 2008; current version published January 30, 2009. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Antonia Papandreou-Suppappola. This work was supported in part by the Israel Science Foundation under Grant 1081/07 and by the European Commission in the framework of the FP7 Network of Excellence in Wireless COMMunications NEWCOM++ (Contract 216715).

The author is with the Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa 32000, Israel (e-mail: yonina@ee.technion.ac.il).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2008.2008212

The James–Stein approach is based on the Stein unbiased risk estimate (SURE) [9], [10], which is an unbiased estimate of the MSE. Since the MSE in general depends on the true unknown parameter values it cannot be used as a design objective. However, using the SURE principle leads to a relatively simple technique for determining methods that have lower MSE than ML. The idea is to choose a class of estimates, and then select the member from the class that minimizes the SURE estimate of the MSE. This strategy has been applied to a variety of different denoising techniques [11]–[14]. Typically, in these problems, implicit prior information on the signal to be recovered is incorporated into the chosen structure of the estimate. For example, in wavelet denoising the signal is assumed to be sparse in the wavelet domain which motivates the use of thresholding. Only the value of the threshold is determined by the SURE principle.

The SURE method is appealing as it allows to directly approximate the MSE of an estimate from the data, without requiring knowledge of the true parameter values. However, it has two main drawbacks which limit its use in practical applications. The first restriction is that it was originally limited to the i.i.d. Gaussian case. Several extensions have been developed for different independent models. In particular, a SURE principle for i.i.d., infinitely divisible random variables with finite variance is derived in [15]. Extensions to independent variables from a continuous exponential family are treated in [16] and [17], while the discrete exponential case is discussed in [18]. All of these generalizations are confined to the independent case which precludes a variety of important applications such as image deblurring.

The second drawback of using SURE as a design criterion is that in order to get meaningful estimators the basic structure of the estimate must be determined in advance. If no parametrization is assumed, then there are too many free variables to be optimized, and the SURE method will typically not lead to good MSE behavior.

In this paper we extend the basic SURE principle in two directions, in order to circumvent the two fundamental drawbacks outlined above. First, we generalize SURE to multivariate, possibly non-i.i.d. exponential families. In particular, we develop an unbiased estimate of the MSE for a general Gaussian vector model. Exponential probability density functions (pdfs) play an important role in statistics due to the Pitman–Koopman–Darmois theorem [19]–[21], which states that among distributions whose domain does not vary with the parameter being estimated, only in exponential families is there a sufficient statistic with bounded dimension as the sample size increases [22]. Furthermore, efficient estimators exist only when the underlying model is exponential. Many known distributions are of

the exponential form, such as Gaussian, gamma, chi-square, beta, Dirichlet, Bernoulli, binomial, multinomial, Poisson, and geometric distributions. Our result has important practical value as it extends the applicability of the SURE technique to more general estimation models, and in particular to scenarios in which the observations are dependent. This is the situation, for example, when using overcomplete wavelet transforms, and in image deblurring. In addition, we derive results for the case in which the model is rank-deficient so that the pdf depends only on a projection of the parameter vector.

An immediate application of this extension is to the general linear Gaussian model. In this setup, we seek to estimate a parameter vector $\boldsymbol{\theta}$ from noisy, blurred measurements $\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$ where \mathbf{w} is a Gaussian noise vector and \mathbf{H} may be rank deficient. One of the most popular recovery strategies in this context is the regularized least-squares (LS) method, in which the estimate is designed to minimize a regularized LS objective. Typical choices of penalization are weighted ℓ_2 or ℓ_1 norms. An important aspect of this technique, which significantly impacts its performance, is selecting the regularization parameter. A variety of different methods have been proposed for this purpose [24]–[31]. When an ℓ_2 norm is used to penalize the LS solution, the resulting estimate is linear and is referred to as Tikhonov regularization. One of the most popular regularization selection methods in this context is generalized cross-validation (GCV) [32]. When an ℓ_1 penalty is used, the resulting estimate is nonlinear so that applying the GCV approach is more complicated. An alternative choice is the discrepancy method in which the parameter is chosen such that the resulting data error is equal to the noise variance.

Here, we suggest a different strategy based on our extended SURE criterion. Specifically, we use SURE to evaluate the MSE of the penalized solution for any choice of regularization, and then select the value that minimizes the SURE estimate. This allows SURE-based optimization of a broad class of deblurring and deconvolution methods including both linear and nonlinear techniques. A similar approach was studied in the special case of Tikhonov regularization with white noise in [24]–[26]. However, our technique is not limited to an ℓ_2 penalty and can be used with any other penalized LS method. When the estimate is not given explicitly but rather as a solution of an optimization problem we can still employ the SURE strategy by using a Monte Carlo approach to approximate the derivative of the estimate, which figures in the SURE expression [39]. Using several test images and a deconvolution problem, we demonstrate that this strategy often leads to significant performance improvement over the standard GCV and discrepancy selection criteria in the context of image deblurring and deconvolution.

Finally, to circumvent the need for predefining a particular structure when applying SURE, we propose an alternative approach based on regularizing the SURE objective. Specifically, we suggest adding a penalization term to the SURE expression and choosing an estimate that minimizes the regularized function. In this way, we can control the properties of the estimate without having to *a priori* assume a specific structure. We then illustrate this strategy in the context of wavelet denoising. Instead of assuming a threshold estimate and choosing the threshold to minimize the SURE criterion, as in [11], we de-

sign an estimate that minimizes an ℓ_1 regularized SURE objective. The resulting denoising scheme has the form of a threshold with a particular form of shrinkage, that is different than that obtained with soft or hard thresholding. To evaluate our method, we compare it with SureShrink of Donoho and Johnstone [11], by repeating the simulations reported in their paper. As we show, the recovery results tend to be better using our technique. Moreover, our approach is general as it is not tailored to a specific problem. We thus believe that using a regularized SURE principle together with the generalized SURE developed here can extend the applicability of SURE-based estimators to a broad class of problems.

The remaining of the paper is organized as follows. In Section II, we introduce the basic concept of MSE estimation. An extension of SURE to exponential families is developed in Section III. In Section IV, we discuss rank-deficient models in which the pdf of the data depends only on a projection of the unknown parameter vector. We then specialize the results to the linear Gaussian model in Section V. Applications to regularization selection are discussed in Section VI. The regularized SURE criterion, together with an application to wavelet denoising, are developed in Section VII.

II. MSE ESTIMATION

We denote vectors by boldface lowercase letters, e.g., \mathbf{x} , and matrices by boldface uppercase letters e.g., \mathbf{A} . The i th component of a vector \mathbf{y} is written as y_i , and $\hat{(\cdot)}$ is an estimated vector. The identity matrix is written as \mathbf{I} , \mathbf{A}^T is the transpose of \mathbf{A} , and \mathbf{A}^\dagger denotes the pseudo-inverse. For a length- m vector function $\mathbf{h}(\mathbf{u}) \in \mathbb{R}^m$ of $\mathbf{u} \in \mathbb{R}^m$,

$$\text{Tr} \left(\frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}} \right) = \sum_{i=1}^m \frac{\partial h_i(\mathbf{u})}{\partial u_i}. \quad (1)$$

We consider the class of problems in which our goal is to estimate a deterministic parameter vector $\boldsymbol{\theta}$ from observations \mathbf{x} which are related through a pdf $f(\mathbf{x}; \boldsymbol{\theta})$. We further assume that the pdf belongs to the exponential family of distributions and can be expressed in the form

$$f(\mathbf{x}; \boldsymbol{\theta}) = r(\mathbf{x}) \exp\{\boldsymbol{\theta}^T \boldsymbol{\phi}(\mathbf{x}) - g(\boldsymbol{\theta})\} \quad (2)$$

where $r(\mathbf{x})$ and $\boldsymbol{\phi}(\mathbf{x})$ are functions of the data only, and $g(\boldsymbol{\theta})$ depends on the unknown parameter $\boldsymbol{\theta}$.

As an example of an application where the model (2) can occur, consider the location problem of estimating a parameter vector $\boldsymbol{\theta} \in \mathbb{R}^m$ from observations $\mathbf{x} \in \mathbb{R}^n$ related through the linear model

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}, \quad (3)$$

where \mathbf{w} is a zero-mean Gaussian random vector with covariance $\mathbf{C} \succ 0$. The pdf of \mathbf{x} is then given by (2) with

$$\begin{aligned} r(\mathbf{x}) &= \frac{1}{\sqrt{(2\pi)^n \det(\mathbf{C})}} \exp \left\{ -\frac{1}{2} \mathbf{x}^T \mathbf{C}^{-1} \mathbf{x} \right\}; \\ \boldsymbol{\phi}(\mathbf{x}) &= \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}; \\ g(\boldsymbol{\theta}) &= \frac{1}{2} \boldsymbol{\theta}^T \mathbf{H}^T \mathbf{C}^{-1} \mathbf{H} \boldsymbol{\theta}. \end{aligned} \quad (4)$$

Other examples of distributions in the exponential family include Poisson with unknown mean, exponential with unknown mean, gamma, and Bernoulli or binomial with unknown success probabilities.

Given the model (2), a sufficient statistic for estimating θ is

$$\mathbf{u} = \phi(\mathbf{x}). \quad (5)$$

Therefore, any reasonable estimate of \mathbf{x} will be a function only of \mathbf{u} . More specifically, from the Rao–Blackwell theorem [33] it follows that if $\hat{\theta}$ is an estimate of θ which is not a function only of \mathbf{u} , then the estimate $E\{\hat{\theta}|\mathbf{u}\}$ has lesser or equal MSE than that of $\hat{\theta}$, for all θ . Therefore, in the sequel, we only consider methods that depend on the data via \mathbf{u} .

Let $\hat{\theta}$ be an arbitrary estimate of θ , which we would like to design to minimize the MSE, defined by $E\{\|\hat{\theta} - \theta\|^2\}$. In practice, $\hat{\theta} = \mathbf{h}(\mathbf{u})$ where $\mathbf{h}(\mathbf{u})$ is some function of \mathbf{u} that is typically chosen to have a particular structure, parameterized by a vector α . For example, $\mathbf{h}(\mathbf{u}) = \alpha\mathbf{u}$ with α a scalar, or $h_i(\mathbf{u}) = \psi_\alpha(u_i)$ where

$$\psi_\alpha(u) = \text{sign}(u)[|u| - \alpha]_+ \quad (6)$$

is a soft-threshold with cut-off α . Ideally, we would like to select α to minimize the MSE. Since this is impossible, as we show below, instead in the SURE approach α is designed to minimize an unbiased estimate (referred to as the SURE estimate) of the MSE.

We can express the MSE of $\hat{\theta} = \mathbf{h}(\mathbf{u})$ as

$$E\{\|\hat{\theta} - \theta\|^2\} = \|\theta\|^2 + E\{\|\mathbf{h}(\mathbf{u})\|^2\} - 2E\{\mathbf{h}^T(\mathbf{u})\theta\}. \quad (7)$$

In order to minimize the MSE over $\mathbf{h}(\mathbf{u})$ we need to explicitly evaluate the expression

$$v(\mathbf{h}, \theta) = E\{\|\mathbf{h}(\mathbf{u})\|^2\} - 2E\{\mathbf{h}^T(\mathbf{u})\theta\}. \quad (8)$$

Evidently, the MSE will depend in general on θ , which is unknown, and therefore cannot be minimized. Instead, we may seek an unbiased estimate of $v(\mathbf{h}, \theta)$ and then choose \mathbf{h} to minimize this estimate. The difficult expression to approximate is $E\{\mathbf{h}^T(\mathbf{u})\theta\}$ as the dependency on θ is explicit. Therefore, we concentrate on estimating this term. If this can be done, then we can easily obtain an unbiased MSE estimate. Specifically, suppose we construct a function $g(\mathbf{h}(\mathbf{u}))$ that depends only on \mathbf{u} (and not on θ), such that

$$E\{g(\mathbf{h}(\mathbf{u}))\} = E\{\mathbf{h}^T(\mathbf{u})\theta\} \triangleq E_{\mathbf{h}, \theta}. \quad (9)$$

Then

$$\hat{v}(\mathbf{h}) = \|\mathbf{h}(\mathbf{u})\|^2 - 2g(\mathbf{h}(\mathbf{u})) \quad (10)$$

is an unbiased estimate of $v(\mathbf{h}, \theta)$, since clearly $E\{\hat{v}(\mathbf{h})\} = v(\mathbf{h}, \theta)$. A reasonable strategy therefore is to select $\mathbf{h}(\mathbf{u})$ to minimize our assessment $\hat{v}(\mathbf{h})$ of the MSE. This approach was first proposed by Stein in [9] and [10] for the i.i.d. Gaussian model (3) with $\mathbf{C} = \mathbf{I}$ and $\mathbf{H} = \mathbf{I}$.

The design framework proposed above reduces to finding an unbiased estimate of $E_{\mathbf{h}, \theta}$. Clearly, any such approximation will depend on the pdf $f(\mathbf{x}; \theta)$. In Section III we develop an unbiased estimate when the pdf is given by the exponential model (2). Before addressing the general setting, to ease the exposition we illustrate the main idea proposed by Stein, by first considering the simpler i.i.d. Gaussian case. In this setting we seek to estimate a vector $\theta \in \mathbb{R}^m$ from measurements $\mathbf{x} = \theta + \mathbf{w}$, where \mathbf{w} is a zero-mean Gaussian random vector with i.i.d. components of variance σ^2 . In Section IV we treat the more difficult case in which \mathbf{u} lies in a subspace \mathcal{A} of \mathbb{R}^m , and the pdf (2) depends on θ only through its orthogonal projection onto \mathcal{A} . This situation arises, for example, in the linear Gaussian model (3) when \mathbf{H} is rank deficient. For this setup, we develop a SURE estimate of the MSE in estimating the projected parameter. In Sections V and VI we consider several examples of estimates in which the free parameters are chosen to minimize the SURE objective. In particular, we propose alternatives to the popular GCV and discrepancy methods for regularization. In Section VII, we suggest a regularized SURE strategy for determining $\mathbf{h}(\mathbf{u})$ without the need for parametrization, and demonstrate its performance in the context of wavelet denoising.

III. EXTENDED SURE PRINCIPLE

A. i.i.d. Gaussian Model

We begin our development by treating the i.i.d. Gaussian setting. Since from (4), $\mathbf{u} = (1/\sigma^2)\mathbf{x}$, we consider estimates $\hat{\theta} = \mathbf{h}(\mathbf{x})$ that are a function of \mathbf{x} .

To develop an unbiased estimate of $E_{\mathbf{h}, \theta}$, we exploit the fact that for the Gaussian pdf $f(\mathbf{x}; \theta)$

$$(x_i - \theta_i)f(\mathbf{x}; \theta) = -\sigma^2 \frac{\partial f(\mathbf{x}; \theta)}{\partial x_i}. \quad (11)$$

Assuming that $E\{|h_i(\mathbf{x})|\}$ is bounded and $h_i(\mathbf{x})$ is weakly differentiable¹ in \mathbf{x} , we have that

$$\begin{aligned} E_{\mathbf{h}, \theta} &= \sum_{i=1}^m \int_{-\infty}^{\infty} h_i(\mathbf{x}) \theta_i f(\mathbf{x}; \theta) d\mathbf{x} \\ &= \sum_{i=1}^m \int_{-\infty}^{\infty} h_i(\mathbf{x}) \left(x_i f(\mathbf{x}; \theta) + \sigma^2 \frac{\partial f(\mathbf{x}; \theta)}{\partial x_i} \right) d\mathbf{x} \\ &= E\{\mathbf{h}^T(\mathbf{x})\mathbf{x}\} + \sigma^2 \sum_{i=1}^m \int_{-\infty}^{\infty} h_i(\mathbf{x}) \frac{\partial f(\mathbf{x}; \theta)}{\partial x_i} d\mathbf{x} \end{aligned} \quad (12)$$

where the second equality is a result of (11). To evaluate the second term in (12), we use integration by parts:

$$\begin{aligned} \int_{-\infty}^{\infty} h_i(\mathbf{x}) \frac{\partial f(\mathbf{x}; \theta)}{\partial x_i} d\mathbf{x} &= - \int_{-\infty}^{\infty} h'_i(\mathbf{x}) f_i(\mathbf{x}; \theta) d\mathbf{x} \\ &= -E\{h'_i(\mathbf{x})\} \end{aligned} \quad (13)$$

¹Roughly speaking, a function is weakly differentiable if it has a derivative almost everywhere, as long as the points that are not differentiable are not delta functions; see [34] for a more formal definition.

were we denoted $h'_i(\mathbf{x}) = \partial h_i(\mathbf{x})/\partial x_i$, and used the fact that $|h_i(\mathbf{x})f(\mathbf{x};\boldsymbol{\theta})| \rightarrow 0$ for $|x_i| \rightarrow \infty$ since $E\{|h_i(\mathbf{x})|\}$ is bounded. We conclude from (12) and (13) that

$$E_{\mathbf{h},\boldsymbol{\theta}} = -\sigma^2 \sum_{i=1}^m E\{h'_i(\mathbf{x})\} + E\{\mathbf{h}^T(\mathbf{x})\mathbf{x}\} \quad (14)$$

and therefore

$$-\sigma^2 \sum_{i=1}^m \frac{\partial h_i(\mathbf{x})}{\partial x_i} + \mathbf{h}^T(\mathbf{x})\mathbf{x} \quad (15)$$

is an unbiased estimate of $E_{\mathbf{h},\boldsymbol{\theta}}$.

B. Extended Sure

We now extend the basic approach outlined in the previous section to the general class of exponential pdfs. In order to address this model, we only consider methods that depend on the data via \mathbf{u} . This enables the use of integration by parts, similar to the i.i.d. Gaussian setting.

Since $\hat{\boldsymbol{\theta}}$ is a function of \mathbf{u} , we can compute the expected value of $\mathbf{h}^T(\mathbf{u})\boldsymbol{\theta}$ using the pdf of \mathbf{u} , rather than that of \mathbf{x} . Combining (2) with the relation $\mathbf{u} = \phi(\mathbf{x})$, it follows that \mathbf{u} also has an exponential pdf

$$f(\mathbf{u};\boldsymbol{\theta}) = q(\mathbf{u}) \exp\{\boldsymbol{\theta}^T \mathbf{u} - g(\boldsymbol{\theta})\}. \quad (16)$$

The pdf of (16) can be derived by utilizing the rules for transformations of variables [33, p. 127].

The following theorem provides an unbiased estimate of $E\{\mathbf{h}^T(\mathbf{u})\boldsymbol{\theta}\}$ which depends only on \mathbf{u} and not on the unknown parameters $\boldsymbol{\theta}$.

Theorem 1: Let \mathbf{x} denote a random vector with exponential pdf given by (2), and let $\mathbf{u} = \phi(\mathbf{x})$ be a sufficient statistic for estimating $\boldsymbol{\theta}$ from \mathbf{x} with pdf (16). Let $\mathbf{h}(\mathbf{u})$ be an arbitrary function of $\boldsymbol{\theta}$ that is weakly differentiable in \mathbf{u} and such that $E\{|h_i(\mathbf{u})|\}$ is bounded. Then

$$E\{\mathbf{h}^T(\mathbf{u})\boldsymbol{\theta}\} = -E\left\{\text{Tr}\left(\frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}}\right)\right\} - E\left\{\mathbf{h}^T(\mathbf{u})\frac{\partial \ln q(\mathbf{u})}{\partial \mathbf{u}}\right\}.$$

From the theorem, it follows that

$$-\text{Tr}\left(\frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}}\right) - \mathbf{h}^T(\mathbf{u})\frac{\partial \ln q(\mathbf{u})}{\partial \mathbf{u}} \quad (17)$$

is an unbiased estimate of $E\{\mathbf{h}^T(\mathbf{u})\boldsymbol{\theta}\}$. In the i.i.d. Gaussian case, $\mathbf{u} = (1/\sigma^2)\mathbf{x}$ so that

$$\frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}} = \sigma^2 \frac{\partial \mathbf{h}(\mathbf{x})}{\partial \mathbf{x}}. \quad (18)$$

Furthermore, \mathbf{u} is a Gaussian i.i.d. vector with elements that have mean $(1/\sigma^2)\boldsymbol{\theta}$ and variance $1/\sigma^2$ so that $q(\mathbf{u}) = K \exp\{-\|\mathbf{u}\|^2 \sigma^2/2\}$ for a constant K . Consequently

$$\frac{\partial \ln q(\mathbf{u})}{\partial \mathbf{u}} = -\sigma^2 \mathbf{u} = -\mathbf{x}. \quad (19)$$

Substituting (18) and (19) into (17), the estimate reduces to (15), derived in the i.i.d. Gaussian setting.

Proof: Using the pdf (16) of \mathbf{u} ,

$$\begin{aligned} E\{\mathbf{h}^T(\mathbf{u})\boldsymbol{\theta}\} &= \int \mathbf{h}^T(\mathbf{u})\boldsymbol{\theta} \exp\{\boldsymbol{\theta}^T \mathbf{u} - g(\boldsymbol{\theta})\} q(\mathbf{u}) d\mathbf{u} \\ &= \sum_{i=1}^m \int h_i(\mathbf{u}) \theta_i \exp\{\boldsymbol{\theta}^T \mathbf{u} - g(\boldsymbol{\theta})\} q(\mathbf{u}) d\mathbf{u}. \end{aligned} \quad (20)$$

Noting that

$$\theta_i \exp\{\boldsymbol{\theta}^T \mathbf{u} - g(\boldsymbol{\theta})\} = \frac{\partial \exp\{\boldsymbol{\theta}^T \mathbf{u} - g(\boldsymbol{\theta})\}}{\partial u_i} \quad (21)$$

we have

$$\begin{aligned} &\int_{-\infty}^{\infty} h_i(\mathbf{u}) \theta_i \exp\{\boldsymbol{\theta}^T \mathbf{u} - g(\boldsymbol{\theta})\} q(\mathbf{u}) d\mathbf{u}_i \\ &= \int_{-\infty}^{\infty} h_i(\mathbf{u}) q(\mathbf{u}) \frac{\partial \exp\{\boldsymbol{\theta}^T \mathbf{u} - g(\boldsymbol{\theta})\}}{\partial u_i} d\mathbf{u}_i \\ &= - \int_{-\infty}^{\infty} \frac{\partial h_i(\mathbf{u}) q(\mathbf{u})}{\partial u_i} \exp\{\boldsymbol{\theta}^T \mathbf{u} - g(\boldsymbol{\theta})\} d\mathbf{u}_i \end{aligned} \quad (22)$$

where we used the fact that $|h_i(\mathbf{u}) q(\mathbf{u}) \exp\{\boldsymbol{\theta}^T \mathbf{u} - g(\boldsymbol{\theta})\}| \rightarrow 0$ for $|u_i| \rightarrow \infty$ since $E\{|h_i(\mathbf{u})|\}$ is bounded. Now,

$$\frac{\partial h_i(\mathbf{u}) q(\mathbf{u})}{\partial u_i} = \frac{\partial h_i(\mathbf{u})}{\partial u_i} q(\mathbf{u}) + \frac{\partial q(\mathbf{u})}{\partial u_i} h_i(\mathbf{u}). \quad (23)$$

Substituting (22) and (23) into (20),

$$\begin{aligned} E\{\mathbf{h}^T(\mathbf{u})\boldsymbol{\theta}\} &= - \sum_{i=1}^m \int \frac{\partial h_i(\mathbf{u}) q(\mathbf{u})}{\partial u_i} \exp\{\boldsymbol{\theta}^T \mathbf{u} - g(\boldsymbol{\theta})\} d\mathbf{u} \\ &= \sum_{i=1}^m \left(-E\left\{\frac{\partial h_i(\mathbf{u})}{\partial u_i}\right\} - E\left\{\frac{\partial q(\mathbf{u})}{\partial u_i} \frac{h_i(\mathbf{u})}{q(\mathbf{u})}\right\} \right) \\ &= -E\left\{\text{Tr}\left(\frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}}\right)\right\} - E\left\{\mathbf{h}^T(\mathbf{u})\frac{\partial \ln q(\mathbf{u})}{\partial \mathbf{u}}\right\} \end{aligned} \quad (24)$$

which completes the proof of the theorem.

Based on Theorem 1 we can develop a generalized SURE principle for estimating an unknown parameter vector $\boldsymbol{\theta}$ in an exponential model. Specifically, let $\hat{\boldsymbol{\theta}} = \mathbf{h}(\mathbf{u})$ be an arbitrary estimate of $\boldsymbol{\theta}$ based on the data \mathbf{x} , where $\mathbf{h}(\mathbf{u})$ satisfies the regularity conditions of Theorem 1. Combining (7) and Theorem 1, an unbiased estimate of the MSE of $\hat{\boldsymbol{\theta}}$ is given by

$$S(\mathbf{h}) = \|\boldsymbol{\theta}\|^2 + \|\mathbf{h}(\mathbf{u})\|^2 + 2\text{Tr}\left(\frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}}\right) + 2\mathbf{h}^T(\mathbf{u})\frac{\partial \ln q(\mathbf{u})}{\partial \mathbf{u}}. \quad (25)$$

We may then design $\hat{\boldsymbol{\theta}}$ by choosing $\mathbf{h}(\mathbf{u})$ to minimize $S(\mathbf{h})$.

In the i.i.d. Gaussian case, (25) reduces to

$$\|\boldsymbol{\theta}\|^2 + \|\mathbf{h}(\mathbf{x})\|^2 + 2\sigma^2 \frac{\partial \mathbf{h}(\mathbf{x})}{\partial \mathbf{x}} - 2\mathbf{h}^T(\mathbf{x})\mathbf{x} \quad (26)$$

where we used the fact that $\mathbf{u} = (1/\sigma^2)\mathbf{x}$ which implies $\partial \ln q(\mathbf{u})/\partial \mathbf{u} = -\mathbf{x}$, and $\partial \mathbf{h}(\mathbf{u})/\partial \mathbf{u} = \sigma^2 \partial \mathbf{h}(\mathbf{x})/\partial \mathbf{x}$. The MSE estimate (26) was first proposed by Stein² in [9] and [10].

The SURE estimate can be used to determine unknown regularization parameters which comprise a given estimation strategy. An example is the SureShrink approach to wavelet denoising [11]. Extending this technique, our general SURE objective (25) may be used to select regularization parameters in more general models. We discuss these ideas in the context of linear Gaussian problems in Section V.

IV. RANK-DEFICIENT MODELS

In some settings, the sufficient statistic \mathbf{u} lies in a subspace \mathcal{A} of \mathbb{R}^m . As an example, suppose that in the Gaussian model (3) \mathbf{H} is rank-deficient. In this case $\mathbf{u} = \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}$ lies in the range space $\mathcal{R}(\mathbf{H}^T)$ of \mathbf{H}^T , which is a subspace of \mathbb{R}^m . If $\boldsymbol{\theta}$ is not restricted to a subspace, then we do not expect to be able to reliably estimate $\boldsymbol{\theta}$ from \mathbf{u} , unless some additional information on $\boldsymbol{\theta}$ is known. Nonetheless, we may still obtain a reliable assessment of the part of $\boldsymbol{\theta}$ that lies in \mathcal{A} . Denote by \mathbf{P} the orthogonal projection onto \mathcal{A} . Then, we show below, that if \mathbf{u} depends on $\boldsymbol{\theta}$ only through $\mathbf{P}\boldsymbol{\theta}$, and in addition \mathbf{u} has an exponential pdf, then we can obtain a SURE estimate of the error in \mathcal{A} , namely $E\{\|\mathbf{P}\hat{\boldsymbol{\theta}} - \mathbf{P}\boldsymbol{\theta}\|^2\}$. If in addition $\hat{\boldsymbol{\theta}}$ lies in \mathcal{A} , then up to a constant, independent of $\boldsymbol{\theta}$, this approximation is also an unbiased estimate of the true MSE $E\{\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2\}$.

To derive the SURE estimate in this case, we first note that if \mathbf{u} lies in \mathcal{A} , then

$$\boldsymbol{\theta}^T \mathbf{u} = (\mathbf{P}\boldsymbol{\theta})^T (\mathbf{P}\mathbf{u}). \quad (27)$$

Suppose that \mathcal{A} has dimension $r < m$. Since $\mathbf{P}\boldsymbol{\theta}$ lies in an r -dimensional space, it can be expressed in terms of r components in an appropriate basis. Denoting by \mathbf{V} an $m \times r$ matrix with orthonormal vectors that span $\mathcal{A} = \mathcal{R}(\mathbf{P})$, the vector $\mathbf{P}\boldsymbol{\theta}$ can be expressed as $\mathbf{P}\boldsymbol{\theta} = \mathbf{V}\boldsymbol{\theta}'$ for an appropriate length- r vector $\boldsymbol{\theta}'$. Similarly, $\mathbf{P}\mathbf{u} = \mathbf{V}\mathbf{u}'$. Therefore, we can write

$$\boldsymbol{\theta}^T \mathbf{u} = \boldsymbol{\theta}'^T \mathbf{u}' \quad (28)$$

where we used the fact that $\mathbf{V}^T \mathbf{V} = \mathbf{I}$. We assume that \mathbf{u}' is a sufficient statistic for $\boldsymbol{\theta}'$ and that $f(\mathbf{u}'; \boldsymbol{\theta}')$ has an exponential pdf:

$$f(\mathbf{u}'; \boldsymbol{\theta}') = q(\mathbf{u}') \exp\{\boldsymbol{\theta}'^T \mathbf{u}' - g(\boldsymbol{\theta}')\}. \quad (29)$$

Under this assumption, we next derive a SURE estimate of the MSE in \mathcal{A} .

Decomposing the squared-error into its components in \mathcal{A} and its orthogonal complement, the MSE can be written as

$$E\{\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2\} = E\{\|\mathbf{P}\hat{\boldsymbol{\theta}} - \mathbf{P}\boldsymbol{\theta}\|^2\} + E\{\|(\mathbf{I} - \mathbf{P})\hat{\boldsymbol{\theta}} - (\mathbf{I} - \mathbf{P})\boldsymbol{\theta}\|^2\}. \quad (30)$$

²We note that the expression obtained by Stein is slightly different since instead of optimizing $\hat{\boldsymbol{\theta}} = \mathbf{h}(\mathbf{x})$, he considered estimates of the form $\hat{\boldsymbol{\theta}} = \mathbf{x} + \mathbf{h}(\mathbf{x})$ and then optimized $\mathbf{h}(\mathbf{x})$. The two expressions differ by a constant, which does not effect the optimization of $\mathbf{h}(\mathbf{x})$.

If $\hat{\boldsymbol{\theta}}$ lies in \mathcal{A} , then $(\mathbf{I} - \mathbf{P})\hat{\boldsymbol{\theta}} = 0$ and the second term is constant, independent of $\hat{\boldsymbol{\theta}}$. Therefore, in this case, to optimize $\hat{\boldsymbol{\theta}}$ it is sufficient to obtain an unbiased estimate of the first term. As we show below, such an assessment can be derived using similar ideas to those in Theorem 1. Even if $\hat{\boldsymbol{\theta}}$ does not lie in \mathcal{A} , the SURE estimate we develop may be used to approximate the first term. Since \mathbf{u} depends only on $\mathbf{P}\boldsymbol{\theta}$, it is reasonable to restrict attention to estimates $\hat{\boldsymbol{\theta}} = \mathbf{h}_\alpha(\mathbf{u})$ where the parameters α are tuned to minimize the MSE in \mathcal{A} $E\{\|\mathbf{P}\hat{\boldsymbol{\theta}} - \mathbf{P}\boldsymbol{\theta}\|^2\}$, subject to any other prior information we may have, such as norm constraints on $\boldsymbol{\theta}$. In such cases we can use a regularized SURE criterion with the SURE objective being the MSE in \mathcal{A} , as we discuss in Section VII.

We now develop a SURE estimate of the MSE $E\{\|\mathbf{P}\hat{\boldsymbol{\theta}} - \mathbf{P}\boldsymbol{\theta}\|^2\}$. To this end we need to find an unbiased estimate of

$$E\{\hat{\boldsymbol{\theta}}^T \mathbf{P}\boldsymbol{\theta}\} = E\{\hat{\boldsymbol{\theta}}^T \mathbf{V}\boldsymbol{\theta}'\} = E\{(\mathbf{V}^T \hat{\boldsymbol{\theta}})^T \boldsymbol{\theta}'\}. \quad (31)$$

Let $\hat{\boldsymbol{\theta}} = \mathbf{h}(\mathbf{u}) = \mathbf{h}(\mathbf{u}')$ (since $\mathbf{u} = \mathbf{V}\mathbf{u}'$, it is clear that $\hat{\boldsymbol{\theta}}$ is a function of \mathbf{u}'). By our assumption, $\boldsymbol{\theta}'$ has an exponential pdf with sufficient statistic \mathbf{u}' . Therefore, we can apply Theorem 1 to $\mathbf{V}^T \mathbf{h}(\mathbf{u})$ for any function $\mathbf{h}(\mathbf{u})$ that obeys the conditions of the theorem, which yields

$$\begin{aligned} E\{\mathbf{h}^T(\mathbf{u})\mathbf{V}\boldsymbol{\theta}'\} &= -E\left\{\text{Tr}\left(\mathbf{V}^T \frac{\partial \mathbf{h}(\mathbf{u}')}{\partial \mathbf{u}'}\right)\right\} - E\left\{\mathbf{h}^T(\mathbf{u}')\mathbf{V} \frac{\partial \ln q(\mathbf{u}')}{\partial \mathbf{u}'}\right\} \\ &= -E\left\{\text{Tr}\left(\mathbf{P} \frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}}\right)\right\} - E\left\{\mathbf{h}^T(\mathbf{u})\mathbf{V} \frac{\partial \ln q(\mathbf{u}')}{\partial \mathbf{u}'}\right\} \end{aligned} \quad (32)$$

where we used the fact that $\mathbf{P} = \mathbf{V}\mathbf{V}^T$ and

$$\frac{\partial \mathbf{h}(\mathbf{u}')}{\partial \mathbf{u}'} = \frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}} \mathbf{V}. \quad (33)$$

To conclude, suppose that \mathbf{u} lies in a subspace \mathcal{A} and is a sufficient statistic for estimating $\mathbf{P}\boldsymbol{\theta}$, with \mathbf{P} denoting the orthogonal projection onto \mathcal{A} , and let $\mathbf{h}(\mathbf{u})$ be an estimate of $\boldsymbol{\theta}$. Then an unbiased estimate of $E\{\|\mathbf{P}\mathbf{h}(\mathbf{u}) - \mathbf{P}\boldsymbol{\theta}\|^2\}$ is

$$\begin{aligned} S(\mathbf{h}) &= \|\mathbf{P}\boldsymbol{\theta}\|^2 + \|\mathbf{P}\mathbf{h}(\mathbf{u})\|^2 \\ &\quad + 2\text{Tr}\left(\mathbf{P} \frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}}\right) + 2\mathbf{h}^T(\mathbf{u})\mathbf{V} \frac{\partial \ln q(\mathbf{u}')}{\partial \mathbf{u}'} \end{aligned} \quad (34)$$

with $\mathbf{u} = \mathbf{V}\mathbf{u}'$ and \mathbf{V} denoting an orthonormal basis for \mathcal{A} . When $\mathcal{A} = \mathbb{R}^m$, $\mathbf{P} = \mathbf{I}$, $\mathbf{V} = \mathbf{I}$ and (34) reduces to (25).

V. LINEAR GAUSSIAN MODEL

We now specialize the SURE principle to the linear Gaussian model (3). We begin by treating the case in which \mathbf{H} is an $n \times m$ matrix with $n \geq m$ and full column rank. We then discuss the rank-deficient scenario.

A. Full-Rank Model

To use Theorem 1 we need to compute the pdf $q(\mathbf{u})$ of \mathbf{u} . Since $\mathbf{u} = \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}$, it is a Gaussian random vector with mean $\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H}\boldsymbol{\theta}$ and covariance $\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H}$. As $q(\mathbf{u})$ is the function

multiplying the exponential in the pdf of \mathbf{u} , it follows from (4) that

$$q(\mathbf{u}) = K \exp \left\{ -\frac{1}{2} \mathbf{u}^T (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{u} \right\} \quad (35)$$

where K is a constant, independent of \mathbf{u} . Therefore,

$$\frac{\partial \ln q(\mathbf{u})}{\partial \mathbf{u}} = -(\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{u} = -\hat{\boldsymbol{\theta}}_{\text{ML}} \quad (36)$$

with $\hat{\boldsymbol{\theta}}_{\text{ML}}$ denoting the ML estimate of $\boldsymbol{\theta}$ given by

$$\hat{\boldsymbol{\theta}}_{\text{ML}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}. \quad (37)$$

From Theorem 1, we then have

$$E \{ \mathbf{h}^T(\mathbf{u}) \boldsymbol{\theta} \} = -E \left\{ \text{Tr} \left(\frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}} \right) - \mathbf{h}^T(\mathbf{u}) \hat{\boldsymbol{\theta}}_{\text{ML}} \right\}. \quad (38)$$

Using (25) and (36) we conclude that

$$S(\mathbf{h}) = \|\boldsymbol{\theta}\|^2 + \|\mathbf{h}(\mathbf{u})\|^2 + 2 \left(\text{Tr} \left(\frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}} \right) - \mathbf{h}^T(\mathbf{u}) \hat{\boldsymbol{\theta}}_{\text{ML}} \right) \quad (39)$$

is an unbiased estimate of the MSE.

B. Rank-Deficient Model

Next, we consider the linear Gaussian model (3) with a rank-deficient \mathbf{H} .

Suppose that \mathbf{H} has a singular value decomposition $\mathbf{H} = \mathbf{U} \Sigma \mathbf{Q}^T$ for some unitary matrices \mathbf{U} and \mathbf{Q} . Let \mathbf{H} have rank r so that Σ is a diagonal $n \times m$ matrix what the first r diagonal elements equal to $\sigma_i > 0$ and the remaining elements equal 0. In this case, \mathbf{V} is equal to the first r columns of \mathbf{Q} and $\boldsymbol{\theta}' = \mathbf{V}^T \boldsymbol{\theta}$. A sufficient statistic for estimating $\boldsymbol{\theta}'$ is $\mathbf{u}' = \mathbf{V}^T \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}$. Clearly, \mathbf{u}' is a Gaussian random vector with mean $\boldsymbol{\mu}' = \mathbf{V}^T \mathbf{H}^T \mathbf{C}^{-1} \mathbf{H} \boldsymbol{\theta}$ and covariance $\mathbf{C}' = \mathbf{V}^T \mathbf{H}^T \mathbf{C}^{-1} \mathbf{H} \mathbf{V}$. Using the SVD of \mathbf{H} we have that

$$\begin{aligned} \boldsymbol{\mu}' &= \Lambda [\mathbf{U}^T \mathbf{C}^{-1} \mathbf{U}]_r \boldsymbol{\theta}', \\ \mathbf{C}' &= \Lambda [\mathbf{U}^T \mathbf{C}^{-1} \mathbf{U}]_r \end{aligned} \quad (40)$$

where Λ is an $r \times r$ diagonal matrix with diagonal elements $\sigma_i^2 > 0$ and $[\mathbf{A}]_r$ is the $r \times r$ top-left principle block of the matrix \mathbf{A} . Since $\mathbf{C} \succ 0$, \mathbf{C}' is invertible. Therefore, $f(\mathbf{u}'; \boldsymbol{\theta}')$ has the form (29) with

$$q(\mathbf{u}') = K \exp \left\{ -\frac{1}{2} \mathbf{u}'^T \mathbf{C}'^{-1} \mathbf{u}' \right\} \quad (41)$$

where K is a constant.

We therefore conclude from (34) that an unbiased estimate of the MSE $E\{\|\mathbf{P}\mathbf{h}(\mathbf{u}) - \mathbf{P}\boldsymbol{\theta}\|^2\}$ is given by

$$\|\mathbf{P}\boldsymbol{\theta}\|^2 + \|\mathbf{P}\mathbf{h}(\mathbf{u})\|^2 + 2 \left(\text{Tr} \left(\mathbf{P} \frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}} \right) - \mathbf{h}^T(\mathbf{u}) \hat{\boldsymbol{\theta}}_{\text{ML}} \right) \quad (42)$$

where $\hat{\boldsymbol{\theta}}_{\text{ML}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^\dagger \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}$ is an ML estimate. Here we used the fact that $\mathbf{V} \mathbf{C}'^{-1} \mathbf{u}' = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^\dagger \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}$.

We summarize our results for the linear Gaussian model in the following proposition.

Proposition 1: Let \mathbf{x} denote measurements of an unknown parameter vector $\boldsymbol{\theta}$ in the linear Gaussian model (3), where \mathbf{w} is a zero-mean Gaussian random vector with covariance $\mathbf{C} \succ 0$. Let $\mathbf{h}(\mathbf{u})$ with $\mathbf{u} = \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}$ be an arbitrary function of $\boldsymbol{\theta}$ that is weakly differentiable in \mathbf{u} and such that $E\{h_i(\mathbf{u})\}$ is bounded, and let \mathbf{P} be an orthogonal projection onto $\mathcal{R}(\mathbf{H}^T)$. Then

$$E \{ \mathbf{h}^T(\mathbf{u}) \mathbf{P} \boldsymbol{\theta} \} = -E \left\{ \text{Tr} \left(\mathbf{P} \frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}} \right) - \mathbf{h}^T(\mathbf{u}) \hat{\boldsymbol{\theta}}_{\text{ML}} \right\}$$

where

$$\hat{\boldsymbol{\theta}}_{\text{ML}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^\dagger \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \quad (43)$$

is an ML estimate of $\boldsymbol{\theta}$. An unbiased estimate of the MSE $E\{\|\mathbf{P}\mathbf{h}(\mathbf{u}) - \mathbf{P}\boldsymbol{\theta}\|^2\}$ is

$$S(\mathbf{h}) = \|\mathbf{P}\boldsymbol{\theta}\|^2 + \|\mathbf{P}\mathbf{h}(\mathbf{u})\|^2 + 2 \left(\text{Tr} \left(\mathbf{P} \frac{\partial \mathbf{h}(\mathbf{u})}{\partial \mathbf{u}} \right) - \mathbf{h}^T(\mathbf{u}) \hat{\boldsymbol{\theta}}_{\text{ML}} \right). \quad (44)$$

C. Examples

To illustrate the use of the SURE principle in the linear Gaussian model, suppose we consider estimates of the form $\hat{\boldsymbol{\theta}} = \mathbf{h}(\mathbf{u}) = \alpha \hat{\boldsymbol{\theta}}_{\text{ML}}$ where $\hat{\boldsymbol{\theta}}_{\text{ML}}$ is given by (43). To choose α we suggest minimizing the SURE unbiased estimate of the MSE given by Proposition 1. Note that in this case $\mathbf{h}(\mathbf{u}) \in \mathcal{R}(\mathbf{H}^T)$ so that $S(\mathbf{h}) + \|(\mathbf{I} - \mathbf{P})\boldsymbol{\theta}\|^2$ is an unbiased estimate of the total MSE $E\{\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2\}$. Therefore, in order to minimize the MSE it suffices to minimize $S(\mathbf{h})$.

For $\mathbf{h}(\mathbf{u}) = \alpha \hat{\boldsymbol{\theta}}_{\text{ML}}$, minimizing $S(\mathbf{h})$ is equivalent to minimizing

$$\alpha^2 \|\hat{\boldsymbol{\theta}}_{\text{ML}}\|^2 + 2 \left(\alpha \text{Tr}((\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^\dagger) - \alpha \|\hat{\boldsymbol{\theta}}_{\text{ML}}\|^2 \right). \quad (45)$$

The optimal choice of α is

$$\alpha = 1 - \frac{\text{Tr}((\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^\dagger)}{\|\hat{\boldsymbol{\theta}}_{\text{ML}}\|^2} \quad (46)$$

resulting in

$$\hat{\boldsymbol{\theta}} = \left(1 - \frac{\text{Tr}((\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^\dagger)}{\|\hat{\boldsymbol{\theta}}_{\text{ML}}\|^2} \right) \hat{\boldsymbol{\theta}}_{\text{ML}}. \quad (47)$$

The estimate of (47) coincides with the balanced blind minimax method proposed in [7, eq. (45)], which was derived based on a minimax framework [8]. Here we see that the same technique results from applying our generalized SURE criterion. A striking feature of this estimate, proved in [7], is that when $\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H}$ is invertible and its effective dimension is larger than 4, it dominates ML for all values of $\boldsymbol{\theta}$ (see Theorem 3 in [7]). This means that its MSE is always lower than that of the ML method, regardless of the true value of $\boldsymbol{\theta}$.

When $\mathbf{H} = \mathbf{I}$ and $\mathbf{C} = \sigma^2 \mathbf{I}$, (47) reduces to

$$\hat{\boldsymbol{\theta}} = \left(1 - \frac{n\sigma^2}{\|\mathbf{x}\|^2} \right) \mathbf{x} \quad (48)$$

which coincides with Stein's estimate [1]. This technique is known to dominate ML for $n \geq 3$.

If in addition we require that $\alpha \geq 0$, then the estimate of (47) becomes

$$\hat{\theta} = \left[1 - \frac{\text{Tr}((\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^\dagger)}{\|\hat{\theta}_{\text{ML}}\|^2} \right]_+ \hat{\theta}_{\text{ML}} \quad (49)$$

where we used the notation

$$[x]_+ = \begin{cases} x, & x \geq 0 \\ 0, & x < 0. \end{cases} \quad (50)$$

The method of (49) is a positive-part version of (47). In the i.i.d. case, it reduces to the positive-part Stein's estimate [35], which is known to dominate the standard Stein approach (48).

Next, consider the case in which $\mathbf{H} = \mathbf{I}$ and $\mathbf{C} = \mathbf{D}$ with $\mathbf{D} = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$ and suppose we seek a diagonal estimate of the form $\hat{\theta}_i = \alpha_i x_i$. Minimizing $S(\mathbf{h})$ of (44) in this case is equivalent to minimizing

$$\sum_{i=1}^n \alpha_i^2 x_i^2 + 2 \sum_{i=1}^n \sigma_i^2 \alpha_i - 2 \sum_{i=1}^n \alpha_i x_i^2 \quad (51)$$

which yields

$$\alpha_i = 1 - \frac{\sigma_i^2}{x_i^2}. \quad (52)$$

Restricting the coefficients α_i to be non-negative leads to the estimate

$$\hat{\theta}_i = \left[1 - \frac{\sigma_i^2}{x_i^2} \right]_+ x_i. \quad (53)$$

In contrast to $\hat{\theta}$ of (47), which dominates the ML method, it can be proven that the estimate of (53) is not dominating. Thus, we see that by allowing for too many free parameters, we impair the performance of the SURE-based estimate. On the other hand, assuming strong structure, as in (47), severely restricts the class of estimators and consequently limits the possible performance advantage which can be obtained. In Section VII, we suggest a regularized SURE strategy in order to overcome this inherent tradeoff between over-parametrization and performance.

VI. APPLICATION TO REGULARIZATION SELECTION

A popular strategy for solving inverse problems of the form (3) is to use regularization techniques in conjunction with a LS objective. Specifically, the estimate $\hat{\theta}$ is chosen to minimize a regularized LS criterion:

$$(\mathbf{x} - \mathbf{H}\hat{\theta})\mathbf{C}^{-1}(\mathbf{x} - \mathbf{H}\hat{\theta}) + \lambda \|\mathbf{L}\hat{\theta}\| \quad (54)$$

where the norm is arbitrary. Here \mathbf{L} is some regularization operator such as the discretization of a first or second order differential operator that accounts for smoothness properties of θ , and λ is the regularization parameter [27], [28]. An important problem in practice is the selection of λ , which strongly effects the recovery performance. One of the most popular approaches

to choosing λ when the estimate is linear (as is the case when a squared- ℓ_2 norm is used in (54)) is the generalized cross-validation (GCV) method [32]. When the estimate takes on a more complicated nonlinear form, a popular selection method is the discrepancy principle [26].

Based on our generalized SURE criterion, we choose λ to minimize the SURE objective (44). As we demonstrate for the cases in which the norm in (54) is the squared- ℓ_2 or ℓ_1 norms, this method can dramatically outperform GCV and the discrepancy technique in practical applications.

A. Image Deblurring

We first consider the case in which the squared- ℓ_2 norm is used in (54). The solution then has the form

$$\hat{\theta} = (\mathbf{Q} + \lambda \mathbf{L}^T \mathbf{L})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \quad (55)$$

where for brevity we denoted

$$\mathbf{Q} = \mathbf{H}^T \mathbf{C}^{-1} \mathbf{H}. \quad (56)$$

The estimate (55) is commonly referred to as Tikhonov regularization [23].

In the GCV method, λ is chosen to minimize

$$G(\lambda) = \frac{1}{\text{Tr}^2(\mathbf{I} - (\mathbf{Q} + \lambda \mathbf{L}^T \mathbf{L})^{-1} \mathbf{Q})} \sum_{i=1}^n (\mathbf{x}_i - [\mathbf{H}\hat{\theta}]_i)^2. \quad (57)$$

The SURE choice follows directly from minimizing (44). In our simulations below, both minimizations were performed by using the `fmincon` function on Matlab.

To demonstrate the performance of the proposed regularization method, we tested it in the context of image deblurring using the HNO deblurring package for Matlab³ based on [36]. We chose several test images, and blurred them using a Gaussian point-spread function of dimension 9 with standard deviation 6 using the function `psfGauss`. We then added zero-mean, Gaussian white noise with variance σ^2 . In Figs. 1 and 2 we compare the deblurred images resulting from using the Tikhonov estimate (55) with $\mathbf{L} = \mathbf{I}$ where the regularization parameter is chosen according to our new SURE criterion (left) and the GCV method (right), for different noise levels.

As can be seen from the figures, our SURE based approach leads to a substantial performance improvement over the standard GCV criterion. This can also be seen in Tables I and II in which we report the resulting MSE values.

B. Deconvolution Example

As another application of the SURE principle, consider the standard deconvolution problem in which a signal $\theta[\ell]$ is convolved by an impulse response $h[\ell]$ and contaminated by additive white Gaussian noise with variance σ^2 . The observations $x[\ell]$ can be written in the form of the linear model (3) where \mathbf{x} is the vector containing the observations $x[\ell]$, θ consists of the input signal $\theta[\ell]$, and \mathbf{H} is a Toeplitz matrix, representing convolution with the impulse response $h[\ell]$.

³The package is available at <http://www2.imm.dtu.dk/~pch/HNO/>.



Fig. 1. Deblurring of Lena using Tikhonov regularization with (left) SURE and (right) GCV choices of regularization and different noise levels: (a), (b) $\sigma = 0.01$; (c), (d) $\sigma = 0.05$; and (e), (f) $\sigma = 0.1$.

To recover $\theta[\ell]$ we may use a penalized LS approach (54) where we assume that the original signal $\theta[\ell]$ is smooth. This can be accounted for by choosing a penalization of the form $\|\mathbf{L}\theta\|_1$ where \mathbf{L} represents a second order derivative operator. The resulting penalized LS estimate can be determined by solving a quadratic optimization problem for fixed λ . In our simulations, we used CVX, a package for specifying and solving convex programs in Matlab [37].

Since the estimate is non-linear, due to the ℓ_1 penalization, we cannot apply the GCV (57). Instead, a popular approach to tune the parameter λ is to use the discrepancy principle in which λ is chosen such that the residual $\|\mathbf{x} - \mathbf{H}\hat{\theta}\|^2$ is equal to the noise level $n\sigma^2$ [25], [26].

To evaluate the performance of the SURE principle in this context, we consider an example from the Regularization Tools [38] for Matlab. All the problems in this toolbox are discretized versions of the Fredholm integral equation of the first kind:

$$g(s) = \int_a^b K(s, t)\theta(t)dt \quad (58)$$

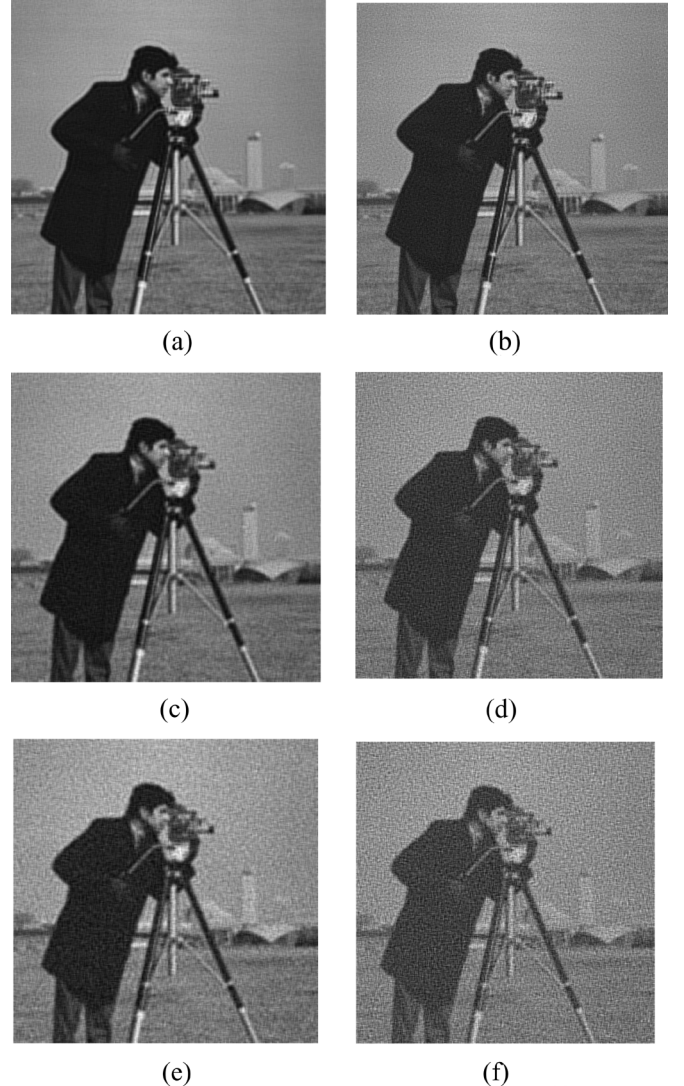


Fig. 2. Deblurring of Cameraman using Tikhonov regularization with SURE (left) and GCV (right) choices of regularization and different noise levels: (a), (b) $\sigma = 0.01$; (c), (d) $\sigma = 0.05$; and (e), (f) $\sigma = 0.1$.

TABLE I
MSE FOR TIKHONOV DEBLURRING OF LENA

	$\sigma = 0.01$	$\sigma = 0.05$	$\sigma = 0.1$
GCV	0.0022	0.0077	0.0133
SURE	0.0011	0.0025	0.0042

TABLE II
MSE FOR TIKHONOV DEBLURRING OF CAMERAMAN

	$\sigma = 0.01$	$\sigma = 0.05$	$\sigma = 0.1$
GCV	0.0033	0.0121	0.0221
SURE	0.0016	0.0039	0.0064

where $K(s, t)$ is the kernel and $\theta(t)$ is the solution for a given $g(s)$. The problem is to estimate $\theta(t)$ from noisy samples of $g(s)$. Using a midpoint rule with n points, (58) reduces to an $n \times n$ linear system $\mathbf{x}_T = \mathbf{H}\theta$. The functions in this toolbox differ in $K(t, s)$ and $\theta(s)$. Below we consider the function heat(n)

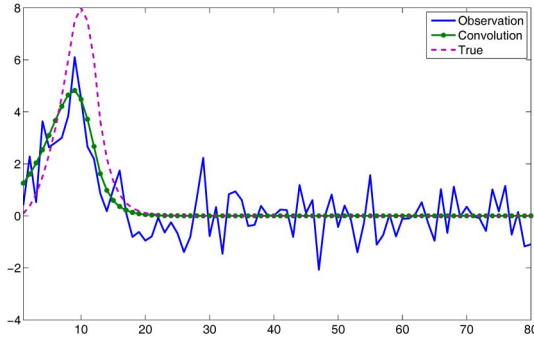


Fig. 3. Original signal θ (dashed), the clean convolved signal (star) and the observations \mathbf{x} with $\sigma = 1$.

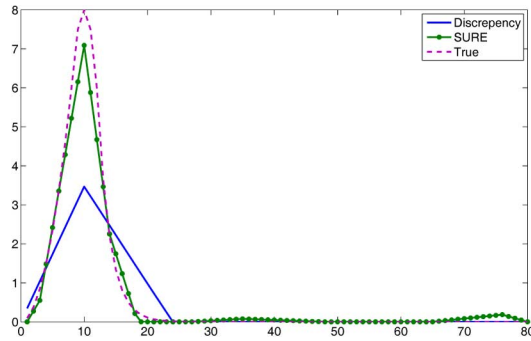


Fig. 4. Deconvolution using weighted ℓ_1 regularization with the discrepancy principle, SURE (star) and the original signal θ (dashed) with $\sigma = 1$.

with $n = 80$. The output of the function is the matrix \mathbf{H} and the true vector θ (which represents $\theta(t)$). The observations are $\mathbf{x} = \mathbf{x}_T + \mathbf{w}$ where \mathbf{w} is a white Gaussian noise vector with variance $\sigma^2 = 1$.

In Fig. 3 we plot the original signal along with the observations \mathbf{x} , and the clean convolved signal $\mathbf{x}_T = \mathbf{H}\theta$. The original signal together with the estimates using the SURE principle and the discrepancy method are plotted in Fig. 4. In order to compute the SURE estimate, we evaluated $\mathbf{h}(\mathbf{u})$ and its derivative numerically for 20 values of λ in the interval $[0.5, 2.5]$. To determine $\mathbf{h}(\mathbf{u})$ we solved the penalized LS problem (54) for the different choices of λ . The gradient of the estimate was computed using the Monte Carlo SURE approach proposed in [39], in which the gradient is approximated by the difference in $\mathbf{h}(\mathbf{u})$ when a small random perturbation is added to \mathbf{u} . We then chose the value of λ on the grid, that minimized the SURE estimate.

It can be seen from Fig. 4 that the SURE method leads to superior performance. The MSE using the SURE approach in this example is 0.10 while the discrepancy strategy leads to an MSE of 1.16.

VII. REGULARIZED SURE METHOD

A crucial element in guaranteeing success of the SURE method is to choose a good parameterization of $\mathbf{h}(\mathbf{u})$. However, in many contexts, such a structure may be hard to find. On the other hand, letting the SURE criterion select many free parameters can deteriorate its performance. One way to

treat this inherent tradeoff is by regularization. Thus, instead of minimizing the SURE objective we suggest minimizing a regularized version:

$$S(\mathbf{h}, \lambda) = S(\mathbf{h}) + \lambda r(\mathbf{h}(\mathbf{u})) \quad (59)$$

where λ is a regularization parameter and $r(\mathbf{h}(\mathbf{u}))$ is a regularization function. For example, we may choose $r(\mathbf{v}) = \|\mathbf{v}\|$ where the norm is arbitrary. The parameter λ is determined by applying the conventional SURE (25) (or (34)) to the estimate $\mathbf{h}(\mathbf{u}, \lambda)$ resulting from solving (59) with λ fixed.

As an example, consider the i.i.d. Gaussian model in which $\mathbf{x} = \theta + \mathbf{w}$ where \mathbf{w} is a Gaussian noise vector with i.i.d. zero-mean components of variance σ^2 . Assuming that θ represents wavelet coefficients of some underlying signal \mathbf{x} , a popular estimation strategy is wavelet denoising in which each component of \mathbf{x} is replaced by a soft or hard-thresholded version. In particular, in their landmark paper, Donoho and Johnstone [11] developed a soft-threshold wavelet denoising method in which

$$\hat{\theta}_i = \begin{cases} |x_i| - t, & |x_i| \geq t \\ 0, & |x_i| < t \end{cases} \quad (60)$$

where t is a threshold value. They suggest selecting t to minimize the SURE criterion, and refer to the resulting estimate as SureShrink (to be more precise, in SureShrink t is determined by SURE only if it lower than some upper limit). In developing the SureShrink method, the function $\mathbf{h}(\mathbf{x})$ is restricted to be a component-wise soft threshold. The motivation for this choice is that the wavelet coefficients below a certain level tend to be sparse. It is well known that soft-thresholding can be obtained as the solution to a LS criterion with an ℓ_1 penalty:

$$\min \{ \|\mathbf{x} - \theta\|^2 + \lambda \|\theta\|_1 \}. \quad (61)$$

Thus, in principle we can view the SureShrink approach as a two-step procedure: We first determine the estimate that minimizes an ℓ_1 penalized LS criterion. We then choose the penalization factor to minimize SURE.

Instead, we suggest choosing an estimate that directly minimizes an ℓ_1 regularized SURE objective, where the only assumption we make is that the processing is performed component wise. Thus, $\hat{\theta}_i = \alpha_i x_i$ for some coefficients $\alpha_i(\mathbf{x}) \geq 0$. Since $\mathbf{u} = (1/\sigma^2)x_i$, $h_i(\mathbf{u}) = \sigma^2 \alpha_i u_i$. With this choice of $\mathbf{h}(\mathbf{u})$, minimizing (59) is equivalent to minimizing the following objective:

$$S(\alpha, \lambda) = \sum_{i=1}^n \alpha_i^2 x_i^2 + 2 \sum_{i=1}^n \alpha_i (\sigma^2 - x_i^2) + \lambda \sum_{i=1}^n |\alpha_i| |x_i|. \quad (62)$$

The optimal choice of $\alpha_i \geq 0$ is

$$\alpha_i = \left[1 - \frac{\sigma^2 + \lambda |x_i|}{x_i^2} \right]_+. \quad (63)$$

The resulting estimate can be viewed as a soft-thresholding method, with a particular choice of shrinkage (different than the standard approach (60)) when the value of x_i exceeds the

TABLE III
MSE FOR DIFFERENT SOFT DENOISING SCHEMES

	Blocks	Bumps	HeaviSine	Doppler
Original	4.054	4.072	4.153	3.945
SureShrink	0.744	0.875	0.205	0.290
RSure	0.694	0.816	0.169	0.273
OracleShrink	0.690	0.828	0.118	0.283

TABLE IV
MSE FOR DIFFERENT DENOISING SCHEMES

	Blocks	Bumps	HeaviSine	Doppler
ScalarShrink	1.043	1.362	0.161	0.594
SteinShrink	1.681	1.730	1.508	1.413

threshold. The threshold is equal to the largest value x_i for which $\alpha_i = 0$ and is given by

$$t = \frac{1}{2} \left(\lambda + \sqrt{\lambda^2 + 4\sigma^2} \right). \quad (64)$$

To choose λ , we substitute the estimate $\hat{\theta}_i = \alpha_i(\lambda)x_i$ with $\alpha_i(\lambda)$ given by (63) into the SURE criterion (44) with $\mathbf{P} = \mathbf{I}$, and minimize with respect to λ . The value of λ can be easily determined numerically.

To demonstrate the advantage of our method over conventional soft-thresholding we implemented our approach on the examples taken from [11]. Specifically, we used the test functions Blocks, Bumps, HeaviSine and Doppler defined in [11]. The length of all signals is 2048 and the noise variance is $\sigma^2 = 4$. We used the Daubechies 8 symmetrical wavelet, and $L = 5$ levels are considered. In Table III we report the empirical MSE values of the original noisy signals, and 3 wavelet denoising schemes: SureShrink which is the method of [11] with the threshold selected using SURE, our proposed regularized SURE method (RSURE), and OracleShrink which is a soft-threshold where the threshold value is selected to minimize the squared-error between the true unknown wavelet coefficient, and its denoised version. Clearly this approach is only for comparison and serves as a benchmark on the best possible performance that can be obtained using any soft threshold. As can be seen from the table, the regularized SURE method performs better in all cases than SureShrink. It is also interesting to see that it sometimes even outperforms OracleShrink which is based on the true unknown θ . The reason the performance can be better than the oracle is that the shrinkage performed in RSURE is different than the conventional soft threshold.

In Table IV we repeat our experiments where now we use the estimates resulting from the standard SURE criterion. Specifically, we consider the positive-part Stein estimate (47) referred to as SteinShrink and the estimate (53) which we refer to as ScalarShrink. Evidently, using the SURE estimate without regularization deteriorates the performance significantly. Thus, SURE alone is not generally sufficient to obtain good estimates. However, adding regularization dramatically improves the behavior without the need to prespecify the desired structure.

TABLE V
MSE FOR DIFFERENT HARD-THRESHOLDING SCHEMES

	Blocks	Bumps	HeaviSine	Doppler
SureShrink	1.902	1.961	0.988	0.630
RSure	1.560	1.912	0.766	0.700

Finally, in Table V we repeat the experiments of Table III to determine the threshold values, but once the values are found we apply hard-thresholding on the coefficients. As can be seen from the table, even though the thresholding operation is now the same in both methods, RSURE performs significantly better. Thus, the threshold determined from this method is superior to that resulting from the SURE criterion without regularization.

VIII. CONCLUSION

In this paper, we developed an unbiased estimate of the MSE in multivariate exponential families by extending the SURE method. This generalized principle can now be used in exponential multivariate estimation problems to develop estimators with improved performance over existing approaches. As an application, we suggested a new strategy for choosing the regularization parameter in penalized inverse problems. We demonstrated via several examples that this method can significantly improve the MSE over the standard GCV and discrepancy approaches. We also suggested a regularized SURE criterion for selecting estimators without the need for prespecifying their structure. Applying this objective in the context of wavelet denoising, we proposed a new type of soft-thresholding which minimizes a penalized estimate of the MSE. As we demonstrated, this strategy can lead to improved MSE behavior in comparison with soft and hard thresholding.

The main contribution of this work is in introducing the generalized SURE criterion and the regularized SURE method and demonstrating their applicability in several examples. In future work, we intend to develop these applications in more detail and further explore the practical use of the proposed design objectives.

ACKNOWLEDGMENT

The author would like to thank Z. Ben-Haim and Prof. M. Elad for many fruitful discussions.

REFERENCES

- [1] C. Stein, "Inadmissibility of the usual estimator for the mean of a multivariate normal distribution," in *Proc. 3rd Berkeley Symp. Math. Statist. Prob.*, Berkeley, CA, 1956, vol. 1, pp. 197–206.
- [2] W. James and C. Stein, "Estimation of quadratic loss," in *Proc. 4th Berkeley Symp. Math. Statist. Prob.*, Berkeley, CA, 1961, vol. 1, pp. 361–379.
- [3] W. E. Strawderman, "Proper Bayes minimax estimators of multivariate normal mean," *Ann. Math. Statist.*, vol. 42, pp. 385–388, 1971.
- [4] K. Alam, "A family of admissible minimax estimators of the mean of a multivariate normal distribution," *Ann. Statist.*, vol. 1, pp. 517–525, 1973.
- [5] J. O. Berger, "Admissible minimax estimation of a multivariate normal mean with arbitrary quadratic loss," *Ann. Statist.*, vol. 4, no. 1, pp. 223–226, Jan. 1976.
- [6] Z. Ben-Haim and Y. C. Eldar, "Blind minimax estimators: Improving on least squares estimation," presented at the IEEE Workshop Statistical Signal Processing (SSP'05), Bordeaux, France, Jul. 2005.

- [7] Z. Ben-Haim and Y. C. Eldar, "Blind minimax estimation," *IEEE Trans. Inf. Theory*, vol. 53, pp. 3145–3157, Sep. 2007.
- [8] Y. C. Eldar, A. Ben-Tal, and A. Nemirovski, "Robust mean-squared error estimation in the presence of model uncertainties," *IEEE Trans. Signal Process.*, vol. 53, no. 1, pp. 168–181, Jan. 2005.
- [9] C. M. Stein, "Estimation of the mean of a multivariate distribution," in *Proc. Prague Symp. Asymptotic Statist.*, 1973, pp. 345–381.
- [10] C. M. Stein, "Estimation of the mean of a multivariate normal distribution," *Ann. Stat.*, vol. 9, no. 6, pp. 1135–1151, Nov. 1981.
- [11] D. L. Donoho and I. M. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *J. Amer. Stat. Assoc.*, vol. 90, no. 432, pp. 1200–1224, Dec. 1995.
- [12] X. P. Zhang and M. D. Desai, "Adapting denoising based on sure risk," *IEEE Signal Process. Lett.*, vol. 5, no. 10, pp. 265–267, 1998.
- [13] F. Luisier, T. Blu, and M. Unser, "A new sure approach to image denoising: Interscale orthonormal wavelet thresholding," *IEEE Trans. Image Process.*, vol. 16, no. 3, pp. 593–606, 2007.
- [14] A. Benazza-Benyahia and J.-C. Pesquet, "Building robust wavelet estimators for multicomponent images using Stein's principle," *IEEE Trans. Image Process.*, vol. 14, no. 11, pp. 1814–1830, 2005.
- [15] R. Averkamp and C. Houdre, "Stein estimation for infinitely divisible laws," *ESAIM: Probabil. Statist.*, pp. 269–269, 2006.
- [16] H. M. Hudson, "A natural identity for exponential families with applications in multiparameter estimation," *Ann. Statist.*, vol. 6, no. 3, pp. 473–484, 1978.
- [17] J. Berger, "Improving on inadmissible estimators in continuous exponential families with applications to simultaneous estimation of gamma scale parameters," *Ann. Statist.*, vol. 8, no. 3, pp. 545–571, 1980.
- [18] J. T. Hwang, "Improving upon standard estimators in discrete exponential families with applications to Poisson and negative binomial cases," *Ann. Statist.*, vol. 10, no. 3, pp. 857–867, 1982.
- [19] E. Pitman, "Sufficient statistics and intrinsic accuracy," *Proc. Camb. Phil. Soc.*, vol. 32, pp. 567–579, 1936.
- [20] G. Darrois, "Sur les Lois de Probabilités à Estimation Exhaustive," *C.R. Acad. Sci. Paris*, vol. 200, pp. 1265–1266, 1935.
- [21] B. Koopman, "On distribution admitting a sufficient statistic," *Trans. Amer. Math. Soc.*, vol. 39, pp. 399–409, 1936.
- [22] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, 2nd ed. New York: Springer-Verlag, 1998.
- [23] A. N. Tikhonov and V. Y. Arsenin, *Solution of Ill-Posed Problems*. Washington, DC: Winston, 1977.
- [24] J. Rice, "Choice of smoothing parameter in deconvolution problems," *Contemp. Math.*, vol. 59, pp. 137–151, 1986.
- [25] L. Desbat and D. Girard, "The 'Minimum reconstruction error' choice of regularization parameters: Some effective methods and their application to deconvolution problem," *SIAM J. Sci. Comput.*, vol. 16, no. 6, pp. 1387–1403, Nov. 1995.
- [26] N. P. Galatsanos and A. K. Katsaggelos, "Methods for choosing the regularization parameter and estimating the noise variance in image restoration and their relation," *IEEE Trans. Image Process.*, vol. 1, no. 3, pp. 322–336, 1992.
- [27] P. C. Hansen, "The use of the L-curve in the regularization of discrete ill-posed problems," *SIAM J. Sci. Stat.ist Comput.*, vol. 14, pp. 1487–1503, 1993.
- [28] M. Hanke and P. C. Hansen, "Regularization methods for large-scale problems," *Surveys Math. Indust.*, vol. 3, no. 4, pp. 253–315, 1993.
- [29] A. Björck, *Numerical Methods for Least-Squares Problems*. Philadelphia, PA: SIAM, 1996.
- [30] R. Molina, A. K. Katsaggelos, and J. Mateos, "Bayesian and regularization methods for hyperparameter estimation in image restoration," *IEEE Trans. Image Process.*, vol. 8, no. 2, pp. 231–246, 1999.
- [31] W. C. Karl, "Regularization in image restoration and reconstruction," in *Handbook of Image and Video Processing*, A. Bovik, Ed., 2nd ed. Amsterdam, The Netherlands: Elsevier, 2005, pp. 183–202.
- [32] G. H. Golub, M. Heath, and G. Wahba, "Generalized cross-validation as a method for choosing a good ridge parameter," *Technometrics*, vol. 21, no. 2, pp. 215–223, May 1979.
- [33] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ: Prentice-Hall, 1993.
- [34] E. H. Lieb and M. Loss, *Analysis*, 2nd ed. Providence, RI: AMS, 2001.
- [35] A. J. Baranchik, Multiple regression and estimation of the mean of a multivariate normal distribution Stanford Univ., Stanford, CA, Tech. Rep. 51, 1964.
- [36] P. C. Hansen, J. G. Nagy, and D. P. O'Leary, *Deblurring Images: Matrices, Spectra, and Filtering*. Philadelphia, PA: SIAM, 2006.
- [37] M. Grant and S. Boyd, CVX: Matlab Software for Disciplined Convex Programming (Web Page and Software) 2008 [Online]. Available: <http://stanford.edu/~boyd/cvx>
- [38] P. C. Hansen, "Regularization tools, a Matlab package for analysis of discrete regularization problems," *Numer. Algorithms*, vol. 6, pp. 1–35, 1994.
- [39] S. Ramani, T. Blu, and M. Unser, "Blind optimization of algorithm parameters for signal denoising by Monte Carlo SURE," presented at the Int. Conf. Acoust., Speech, Signal Processing (ICASSP), Las Vegas, NV, Apr. 2008.



Yonina C. Eldar (S'98–M'02–SM'07) received the B.Sc. degree in physics and the B.Sc. degree in electrical engineering, both from Tel-Aviv University (TAU), Tel-Aviv, Israel, in 1995 and 1996, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 2001.

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 signal processing, 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 to 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*, and the *SIAM Journal on Matrix Analysis and Applications*, and on the Editorial Board of *Foundations and Trends in Signal Processing*.