Unbiased Risk Estimates for Singular Value Thresholding and Spectral Estimators

Emmanuel J. Candès* Carlos A. Sing-Long † Joshua D. Trzasko ‡

October, 2012

Abstract

In an increasing number of applications, it is of interest to recover an approximately low-rank data matrix from noisy observations. This paper develops an unbiased risk estimate—holding in a Gaussian model—for any *spectral estimator* obeying some mild regularity assumptions. In particular, we give an unbiased risk estimate formula for singular value thresholding (SVT), a popular estimation strategy which applies a soft-thresholding rule to the singular values of the noisy observations. Among other things, our formulas offer a principled and automated way of selecting regularization parameters in a variety of problems. In particular, we demonstrate the utility of the unbiased risk estimation for SVT-based denoising of real clinical cardiac MRI series data. We also give new results concerning the differentiability of certain matrix-valued functions.

Keywords. Singular value thresholding, Stein's unbiased risk estimate (SURE), differentiability of eigenvalues and eigenvectors, magnetic resonance cardiac imaging.

1 Introduction

Suppose we have noisy observations $\boldsymbol{Y}$ about an $m \times n$ data matrix $\boldsymbol{X}^0$ of interest,

$$Y_{ij} = X^0_{ij} + W_{ij}, \qquad W_{ij} \overset{\text{iid}}{\sim} \mathcal{N}(0, \tau^2), \quad \begin{array}{l} i = 1, \dots, m, \\ j = 1, \dots, n. \end{array} \tag{1.1}$$

We wish to estimate $\boldsymbol{X}^0$ as accurately as possible. In this paper, we are concerned with situations where the estimand has some structure, namely, $\boldsymbol{X}^0$ has low rank or is well approximated by a low-rank matrix. This assumption is often met in practice since the columns of $\boldsymbol{X}^0$ can be quite correlated. For instance, these columns may be individual frames in a video sequence, which are typically highly correlated. Another example concerns the acquisition of hyperspectral images in which each column of $\boldsymbol{X}^0$ is a 2D image at a given wavelength. In such settings, images at nearby wavelengths typically exhibit strong correlations. Hence, the special low-rank regression problem (1.1) occurs in very many applications and is the object of numerous recent studies.

Recently, promoting low-rank has been identified as a promising tool for denoising series of MR images, such as those that arise in functional MRI (fMRI) [34], relaxometry [22], cardiac MRI [36], NMR spectroscopy [5, 25], and diffusion-weighted imaging [15, 24], among others. In dynamic

*Departments of Statistics and of Mathematics, Stanford University, Stanford, CA 94305
†Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305
‡Department of Physiology and Biomedical Engineering, Mayo Clinic, Rochester, MN 55905

applications like cine cardiac imaging, where "movies" of the beating heart are created, neighboring tissues tend to exhibit similar motion profiles through time due to their physical connectedness. The diversity of temporal behaviors in this settings will, by nature, be limited, and the Casorati matrix (a matrix whose columns comprise vectorized frames of the image series) formed from this data will be low-rank [17]. For example, in breath-hold cine cardiac imaging, background tissue is essentially static and the large submatrix of the Casoratian corresponding to this region is very well approximated by a rank-1 matrix.

1.1 Singular value thresholding

Whenever the object of interest has (approximately) low rank, it is possible to improve upon the naive estimate $\hat{\boldsymbol{X}}_0 = \boldsymbol{Y}$ by regularizing the maximum likelihood. A natural approach consists in truncating the singular value decomposition of the observed matrix $\boldsymbol{Y}$, and solve

$$\mathrm{SVHT}_\lambda(\boldsymbol{Y}) = \arg\min_{\boldsymbol{X}\in\mathbb{R}^{m\times n}} \frac{1}{2}\|\boldsymbol{Y}-\boldsymbol{X}\|_F^2 + \lambda\ \mathrm{rank}(\boldsymbol{X}), \tag{1.2}$$

where λ a positive scalar. As is well known, if

$$\boldsymbol{Y} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^* = \sum_{i=1}^{\min(m,n)} \sigma_i \boldsymbol{u}_i \boldsymbol{v}_i^* \tag{1.3}$$

is a singular value decomposition for $\boldsymbol{Y}$, the solution is given by retaining only the part of the expansion with singular values exceeding λ,

$$\mathrm{SVHT}_\lambda(\boldsymbol{Y}) = \sum_{i=1}^{\min(m,n)} \mathbb{I}(\sigma_i > \lambda)\, \boldsymbol{u}_i \boldsymbol{v}_i^*.$$

In other words, one applies a hard-thresholding rule to the singular values of the observed matrix $\boldsymbol{Y}$. Such an estimator is discontinuous in $\boldsymbol{Y}$ and a popular alternative approach applies, instead, a soft-thresholding rule to the singular values:

$$\mathrm{SVT}_\lambda(\boldsymbol{Y}) = \sum_{i=1}^{\min(m,n)} (\sigma_i - \lambda)_+ \boldsymbol{u}_i \boldsymbol{v}_i^*; \tag{1.4}$$

that is, we shrink the singular values towards zero by a constant amount λ. Here, the estimate $\mathrm{SVT}_\lambda(\boldsymbol{Y})$ is Lipschitz continuous. This follows from the fact that the singular value thresholding operation (1.4) is the prox of the nuclear norm $\|\cdot\|_*$ (the nuclear norm of a matrix is sum of its singular values), i.e. is the unique solution to

$$\min \quad \frac{1}{2}\|\boldsymbol{Y}-\boldsymbol{X}\|_F^2 + \lambda\|\boldsymbol{X}\|_*. \tag{1.5}$$

The focus of this paper is on this smoother estimator referred to as the singular value thresholding (SVT) estimate.

1.2 A SURE formula

The classical question is, of course, how much shrinkage should be applied. Too much shrinkage results in a large bias while too little results in a high variance. To find the correct trade-off, it would be desirable to have a method that would allow us to compare the quality of estimation for different values of the parameter λ. Ideally, we would like to select λ as to minimize the mean-squared error or risk

$$\mathrm{MSE}(\lambda) = \mathbb{E}\,\|\boldsymbol{X}^0 - \mathrm{SVT}_\lambda(\boldsymbol{Y})\|_F^2. \tag{1.6}$$

Unfortunately, this cannot be achieved since the expectation in (1.6) depends on the true $\boldsymbol{X}^0$, and is thus unknown. Luckily, when the observations follow the model (1.1), it is possible to construct an *unbiased* estimate of the risk, namely, Stein's Unbiased Risk Estimate (SURE) [32] given by

$$\mathrm{SURE}(\mathrm{SVT}_\lambda)(\boldsymbol{Y}) = -mn\tau^2 + \sum_{i=1}^{\min(m,n)} \min(\lambda^2, \sigma_i^2) + 2\tau^2 \mathrm{div}\,(\mathrm{SVT}_\lambda(\boldsymbol{Y})), \tag{1.7}$$

where $\{\sigma_i\}_{i=1}^n$ denote the singular values of $\boldsymbol{Y}$. Above, 'div' is the divergence of the nonlinear mapping SVT_λ, which is to be interpreted in a weak sense. Roughly speaking, it can fail to exist on negligible sets. The main contribution of this paper is to provide a closed-form expression for the divergence of this estimator. We prove that in the real-valued case,

$$\mathrm{div}\,(\mathrm{SVT}_\lambda(\boldsymbol{Y})) = \sum_{i=1}^{\min(m,n)} \left[\mathbb{I}(\sigma_i > \lambda) + |m-n|\left(1 - \frac{\lambda}{\sigma_i}\right)_+\right] + 2\sum_{i\neq j, i,j=1}^{\min(m,n)} \frac{\sigma_i(\sigma_i - \lambda)_+}{\sigma_i^2 - \sigma_j^2}, \tag{1.8}$$

when $\boldsymbol{Y}$ is *simple*—i.e., has no repeated singular values—and 0 otherwise, say, is a valid expression for the weak divergence. Hence, this simple formula can be used in (1.7), and ultimately leads to the determination of a suitable threshold level by minimizing the *estimate* of the risk, which only depends upon the observed data.

We pause to present some numerical experiments to assess the quality of SURE as an estimate of the risk. We work with four matrices $\boldsymbol{X}_i^0$, $i = 1, \ldots, 4$, of size 200×500. Here, $\boldsymbol{X}_1^0$ has full rank; $\boldsymbol{X}_2^0$ has rank 100; $\boldsymbol{X}_3^0$ has rank 10; and $\boldsymbol{X}_4^0$ has singular values equal to $\sigma_i = \sqrt{200}/(1+\mathrm{e}^{(i-100)/20})$, $i = 1, \ldots, 200$. Finally, each matrix is normalized so that $\|\boldsymbol{X}_i^0\|_F = 1$, $i = 1, \ldots, 4$. Next, two methods are used to estimate the risk (1.6) of SVT_λ seen as a function of λ. The first methods uses

$$\hat{R}_i(\lambda) = \frac{1}{50}\sum_{j=1}^{50} \|\mathrm{SVT}_\lambda(\boldsymbol{Y}_j^{(i)}) - \boldsymbol{X}_i^0\|_F^2,$$

where $\{\boldsymbol{Y}_i^{(j)}\}_{j=1}^{50}$ are independent samples drawn from model (1.1) with $\boldsymbol{X}^0 = \boldsymbol{X}_i^0$. We set this to be the value of reference. The second uses $\mathrm{SURE}(\mathrm{SVT}_\lambda)(\boldsymbol{Y})$, where $\boldsymbol{Y}$ is drawn from model (1.1) independently from $\{\boldsymbol{Y}_j^{(i)}\}$. Finally, in each case we work with values of the signal-to-noise ratio, defined here as $\mathrm{SNR} = \|\boldsymbol{X}^0\|_F/\sqrt{mn}\tau = 1/\sqrt{mn}\tau$, and set $\mathrm{SNR} = 0.5, 1, 2$ and 4. The results are shown in Figure 1, where one can see that SURE remains very close to the true value of the risk, even though it is calculated from a single observation. Matlab code reproducing the figures is available and computing SURE formulas for various spectral estimators is available at `http://www-stat.stanford.edu/~candes/SURE_SVT_code.zip`.

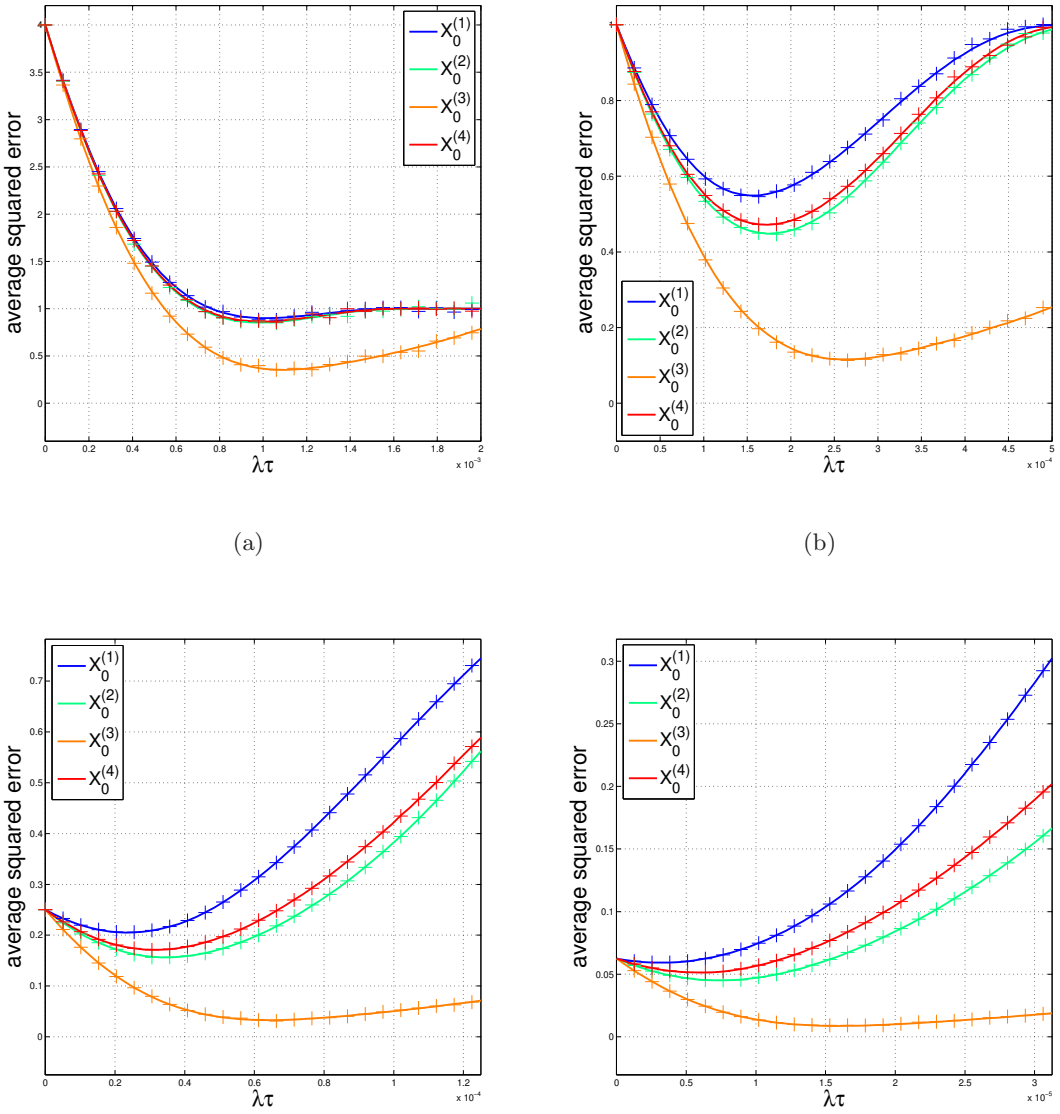

Figure 1: Comparison of the risk estimate using Monte Carlo (solid line) and SURE (cross) versus $\lambda \times \tau$ for $\boldsymbol{X}^0 \in \mathbb{R}^{200\times 500}$ and: (a) SNR $= 0.5$; (b) SNR $= 1$; (c) SNR $= 2$; and (d) SNR $= 4$. The colors indicate the matrices used in each case. As we can see, SURE follows closely the Monte Carlo estimates, even though it requires only one observation to estimate the risk.

In MR applications, observations can take on complex values. Indeed, MRI scanners employ a process known as quadrature detection (two orthogonal phase-sensitive detectors) to observe both the magnitude and phase of rotating magnetization that is induced by radio-frequency (RF) excitation [18]. Quadrature detection both increases SNR and allows for the encoding of motion information like flow. MRI noise, which is Gaussian, is also complex valued. In general, MRI noise can also be assumed iid, noting that inter-channel correlation in parallel MRI data can be removed via Cholesky pre-whitening. Thus, model (1.1) has to be modified as

$$\boldsymbol{Y} = \boldsymbol{X}^0 + \boldsymbol{W}, \qquad \mathrm{Re}(W_{ij}),\, \mathrm{Im}(W_{ij}) \overset{\text{iid}}{\sim} \mathcal{N}(0, \tau^2), \tag{1.9}$$

where the real and imaginary parts are also independent. In this case, SURE becomes

$$\mathrm{SURE}(\mathrm{SVT}_\lambda)(\boldsymbol{Y}) = -2mn\tau^2 + \sum_{i=1}^{\min(m,n)} \min(\lambda^2, \sigma_i^2) + 2\tau^2 \mathrm{div}\left(\mathrm{SVT}_\lambda(\boldsymbol{Y})\right).$$

We also provide an expression for the weak divergence in this context, namely,

$$\mathrm{div}\left(\mathrm{SVT}_\lambda(\boldsymbol{Y})\right) = \sum_{i=1}^{\min(m,n)} \left[\mathbb{I}(\sigma_i > \lambda) + (2|m-n|+1)\left(1 - \frac{\lambda}{\sigma_i}\right)_+\right] + 4 \sum_{i\neq j,\, i,j=1}^{\min(m,n)} \frac{\sigma_i(\sigma_i - \lambda)_+}{\sigma_i^2 - \sigma_j^2}, \tag{1.10}$$

when $\boldsymbol{Y}$ is simple, and 0 otherwise. This formula can be readily applied to MR data as we shall see in Section 2. The decomposition into real and imaginary parts would suggest that the divergence would be proportional to twice the divergence in the real case. However, this is not the case. The most significant difference is that there is a contribution of the inverse of the singular values even when the matrix is square.

1.3 Extensions

The formulae (1.8)–(1.10) have applications beyond SURE. For instance, the divergence of an estimation procedure arises naturally when trying to estimate the degrees of freedom associated with it. Consider the estimate $\mathrm{SVT}_\lambda(\boldsymbol{Y})$ of $\boldsymbol{X}^0$ when the observation $\boldsymbol{Y}$ comes from an additive error model. The degrees of freedom [11, 32] are defined as

$$\mathrm{df}(\mathrm{SVT}_\lambda) = \frac{1}{\tau^2} \sum_{i=1}^{m}\sum_{j=1}^{n} \mathbf{cov}\left(\mathrm{SVT}_\lambda(\boldsymbol{Y})_{ij}, Y_{ij}\right).$$

When the observations $\boldsymbol{Y}$ follow model (1.1), it is possible to write the degrees of freedom as

$$\mathrm{df}(\mathrm{SVT}_\lambda) = \mathbb{E}\{\mathrm{div}(\mathrm{SVT}_\lambda(\boldsymbol{Y}))\}.$$

Therefore, the expression we provide is also useful to estimate or calculate the degrees of freedom of singular value thresholding.

Finally, although our work focuses on SVT, our methods extend beyond this particular estimator. Consider estimators given by *spectral functions*. These act on the singular values and take the form

$$f(\boldsymbol{Y}) = \sum_{i=1}^{\min(m,n)} f_i(\sigma_i)\boldsymbol{u}_i\boldsymbol{v}_i^* := \boldsymbol{U} f(\boldsymbol{\Sigma})\boldsymbol{V}^*, \quad \text{for all } \boldsymbol{Y} \in \mathbb{R}^{m\times n}, \tag{1.11}$$

where $\boldsymbol{Y} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^*$ is any SVD (SVT is in this class). Our methods show that these functions admit a SURE formula, given by

$$\mathrm{SURE}(f)(\boldsymbol{Y}) = -mn\tau^2 + \|f(\boldsymbol{Y}) - \boldsymbol{Y}\|_F^2 + 2\tau^2 \mathrm{div}\,(f(\boldsymbol{Y})),$$

and that under mild assumptions there exists a closed-form for their divergence:

$$\mathrm{div}\,(f(\boldsymbol{X})) = \sum_{i=1}^{\min(m,n)} \left(f_i'(\sigma_i) + |m-n| \frac{f_i'(\sigma_i)}{\sigma_i} \right) + 2 \sum_{i \neq j,\, i,j=1}^{\min(m,n)} \frac{\sigma_i f_i(\sigma_i)}{\sigma_i^2 - \sigma_j^2}. \tag{1.12}$$

This is of interest because such estimators arise naturally in regularized regression problems. For instance, let $J : \mathbb{R}^{m\times n} \mapsto \mathbb{R}$ be a lower semi-continuous, proper convex function of the form

$$J(\boldsymbol{X}) = \sum_{i=1}^{\min(m,n)} J_i(\sigma_i(\boldsymbol{X})),$$

Then, for $\lambda > 0$ the estimator

$$f_\lambda(\boldsymbol{Y}) = \arg\min_{\boldsymbol{X}\in\mathbb{R}^{m\times n}} \frac{1}{2}\|\boldsymbol{Y} - \boldsymbol{X}\|_F^2 + \lambda J(\boldsymbol{X}) \tag{1.13}$$

is spectral. Hence, (1.13) can be used broadly.

1.4 Methods and content

The methods we use to derive our formulas are standard since they involve computing the Jacobian of the SVD, see [10] and [27].[1] During the preparation of this manuscript, the arXiv paper [8] came to our attention. This work does not give a SURE formula but uses the same standard methods to tackle the computation of the Jacobian of the SVT estimate in the real-valued case when the input matrix does not have repeated singular values. Another pointer to this work is in Section 4.

The structure of the paper is as follows. In Section 2 we present applications in MRI illustrating the advantages of choosing the threshold in a disciplined fashion. Section 3 provides precise statements, and a rigorous justification for (1.8) and (1.10). Section 4 deals with the differentiability of spectral functions (1.11), and thus supports Section 3. We conclude with a short discussion of our results, and potential research directions in Section 5.

2 Applications in Magnetic Resonance Imaging (MRI)

SVT is a computationally-straightforward, yet, powerful denoising strategy for MRI applications where spatio-temporal/parameteric behavior is either *a priori* unknown or else defined accordingly to a complicated nonlinear model that may be numerically challenging to work with. A practical challenge in most rank-based estimation problems in MRI (and inverse problems in general), however, lies in the selection of the regularization or threshold parameter. Defining an automated, disciplined, and consistent methodology for selecting this parameter inherently improves clinical

[1] In passing, Iain Johnstone has informed us that a variation on these methods has been used in unpublished work by our colleague Charles Stein.

workflow both by accelerating the tuning process through the use of optimization-based, rather than heuristic, search strategies and by freeing the MRI scanner technician so that they can focus on other patient-specific tasks. Moreover, eliminating the human element from the denoising process mitigates inter- and intra-operator variability, and thus raises diagnostic confidence in the denoising results since images are wholly reproducible.

2.1 Noise reduction in cardiac MRI series

Here, we demonstrate a potential use of the SVT unbiased risk estimate for automated and optimized denoising of dynamic cardiac MRI series. Dynamic cardiac imaging is performed either in cine or real-time mode. Cine MRI, the clinical gold-standard for measuring cardiac function/volumetrics [3], produces a movie of roughly 20 cardiac phases over a single cardiac cycle (heart beat). However, by exploiting the semi-periodic nature of cardiac motion, it is actually formed over many heart beats. Cine sampling is gated to a patient's heart beat, and as each data measurement is captured it is associated with a particular cardiac phase. This process continues until enough data has been collected such that all image frames are complete. Typically, an entire 2D cine cardiac MRI series is acquired within a single breath hold (less than 30 secs).

In real-time cardiac MRI, an image series covering many heart beats is generated. Although the perceived temporal resolution of this series is coarser than that of a cine series, the "temporal footprint" of each frame is actually shorter since it is only comprised of data from one cardiac cycle. Real-time cardiac MRI is of increasing clinical interest for studying transient functional processes such as first-pass myocardial perfusion, where the hemodynamics of an intravenously-injected contrast bolus are visualized. First-pass myocardial perfusion proffers visualization of both damage to heart muscle as well as coronary artery blockage. Typically, a real-time cardiac MRI study will exceed feasible breath-hold time and respiratory motion may be visible.

Simultaneously achieving high spatial resolution and SNR is challenging in both cine and real-time cardiac MRI due to the dynamic nature of the target signal. Signal averaging (NEX ≥ 1), a standard MRI techniques for noise reduction, is infeasible in real-time imaging and undesirable in cine imaging due to potential misregistration artifacts, since it cannot be executed within a single breath-hold. Real-time acquisitions, which are generally based on gradient recalled echo (GRE) protocols, have inherently low SNR due to their use of very short repetition (TR) and echo times (TE). Cine acquisitions use either standard GRE or balanced steady-state free precession (bSSFP) protocols. When imaging with 1.5 T (Tesla) magnets, bSSFP cine sequences can often yield sufficient SNR. However, unlike many other MRI protocols, SSFP sequences do not trivially gain SNR when moved to higher-field systems (≥ 3.0 T) [26] which is an emerging clinical trend. As magnetic field strength is increased, the RF excitation flip angle used in a bSSFP sequence must be lowered to adhere to RF power deposition (SAR) safety limits. This results in weaker signal excitation which can mitigate gains in bulk magnetization. Poor receiver coil sensitivity at the heart's medial location and signal loss due to iron overload (hemochromatosis), among other factors, can further reduce SNR in both imaging strategies. Beyond confounding visual radiological evaluation, even moderate noise levels in a cardiac image can also degrade the performance of automated segmentation methods that are used for quantitative cardiac function evaluation. Therefore, effective denoising techniques that preserve the morphological and dynamic profiles of cardiac image series are clinical valuable.

Consider a series of t separate $n \times n$ 2D MR images. To utilize SVT to denoise this data, it must first be transformed into a $n^2 \times t$ Casorati matrix, $\boldsymbol{Y} = \boldsymbol{X}^0 + \boldsymbol{W}$. In many MRI scenarios, spatial

dimensionality will greatly exceed temporal/parametric dimensionality ($n^2 \gg t$) and $\boldsymbol{Y}$ typically is a very thin matrix. Due to limited degrees-of-freedom, even optimally-parameterized SVT may result in temporal/parametric blurring. One solution to this problem is to analyze the image series in a spatial block-wise manner (see Figure 2) rather than globally [36].

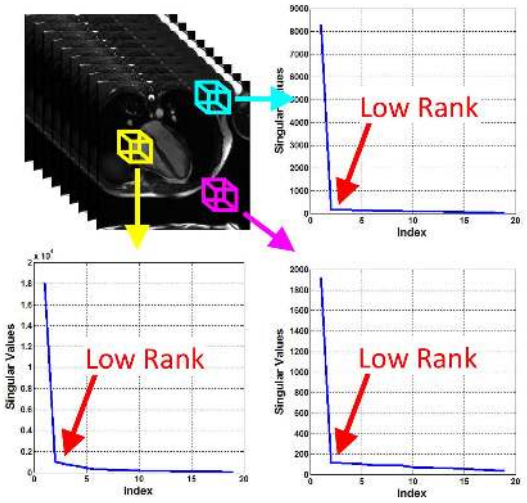

Figure 2: Singular values of the Casorati matrices formed from three different $8 \times 8 \times 19$ block sets extracted from a cine cardiac sequence, including: a static region (cyan); a dynamic region (yellow); and background noise (magenta).

Let $\boldsymbol{R}_b$ be a binary operator that extracts k^2 rows from a matrix corresponding to a $k \times k$ spatial block, specified by index b, within each image. Block-wise SVT can be defined as

$$\mathrm{BSVT}_\lambda(\boldsymbol{Y}) = c^{-1} \sum_{b \in \Omega} \boldsymbol{R}_b^* \, \mathrm{SVT}_\lambda(\boldsymbol{R}_b \boldsymbol{Y}), \tag{2.1}$$

where Ω denotes a set of (potentially overlapping) blocks that uniformly tiles the image domain, i.e., $\sum_{b \in \Omega} \boldsymbol{R}_b^* \boldsymbol{R}_b = c \boldsymbol{I}_{n^2 \times n^2}$, $c > 0$. In words, (2.1) performs SVT on a family of submatrices of $\boldsymbol{Y}$ and accumulates a weighted sum of the results. Of course, for $k = n$ and $|\Omega| = 1$, (2.1) resorts to standard SVT. The unbiased risk estimator developed in the previous section readily extends for this generalized SVT model. By linearity,

$$\mathrm{div}\, \mathrm{BSVT}_\lambda(\boldsymbol{Y}) = c^{-1} \sum_{b \in \Omega} \mathrm{div}\, \boldsymbol{R}_b^* \, \mathrm{SVT}_\lambda(\boldsymbol{R}_b \boldsymbol{Y}). \tag{2.2}$$

Extending the identity (10) from [2] for matrices then asserts

$$\mathrm{div}\, \boldsymbol{R}_b^* \, \mathrm{SVT}_\lambda(\boldsymbol{R}_b \boldsymbol{Y}) = \sum_{ij} \frac{\partial (\boldsymbol{R}_b^* \, \mathrm{SVT}_\lambda(\boldsymbol{R}_b \boldsymbol{Y}))_{ij}}{\partial Y_{ij}} = \sum_{ij} \frac{\partial \, \mathrm{SVT}_\lambda(\boldsymbol{R}_b \boldsymbol{Y})_{ij}}{\partial (\boldsymbol{R}_b \boldsymbol{Y})_{ij}} (\boldsymbol{R}_b \boldsymbol{R}_b^*)_{ii}. \tag{2.3}$$

Now, observing that $\boldsymbol{R}_b \boldsymbol{R}_b^* = \boldsymbol{I}_{k^2 \times k^2}$, $\forall b \in \Omega$, we see that $\mathrm{div}\, \boldsymbol{R}_b^* \, \mathrm{SVT}_\lambda(\boldsymbol{R}_b \boldsymbol{Y}) = \mathrm{div}\, \mathrm{SVT}_\lambda(\boldsymbol{R}_b \boldsymbol{Y})$, where the latter term is as in (1.8) for $\boldsymbol{Y} \in \mathbb{R}^{n^2 \times t}$ or (1.10) for $\boldsymbol{Y} \in \mathbb{C}^{n^2 \times t}$. Define the mean-squared error of BSVT_λ as

$$\mathrm{MSE}(\lambda) = \mathbb{E} \, \| \boldsymbol{X}_0 - \mathrm{BSVT}_\lambda(\boldsymbol{Y}) \|_F^2, \tag{2.4}$$

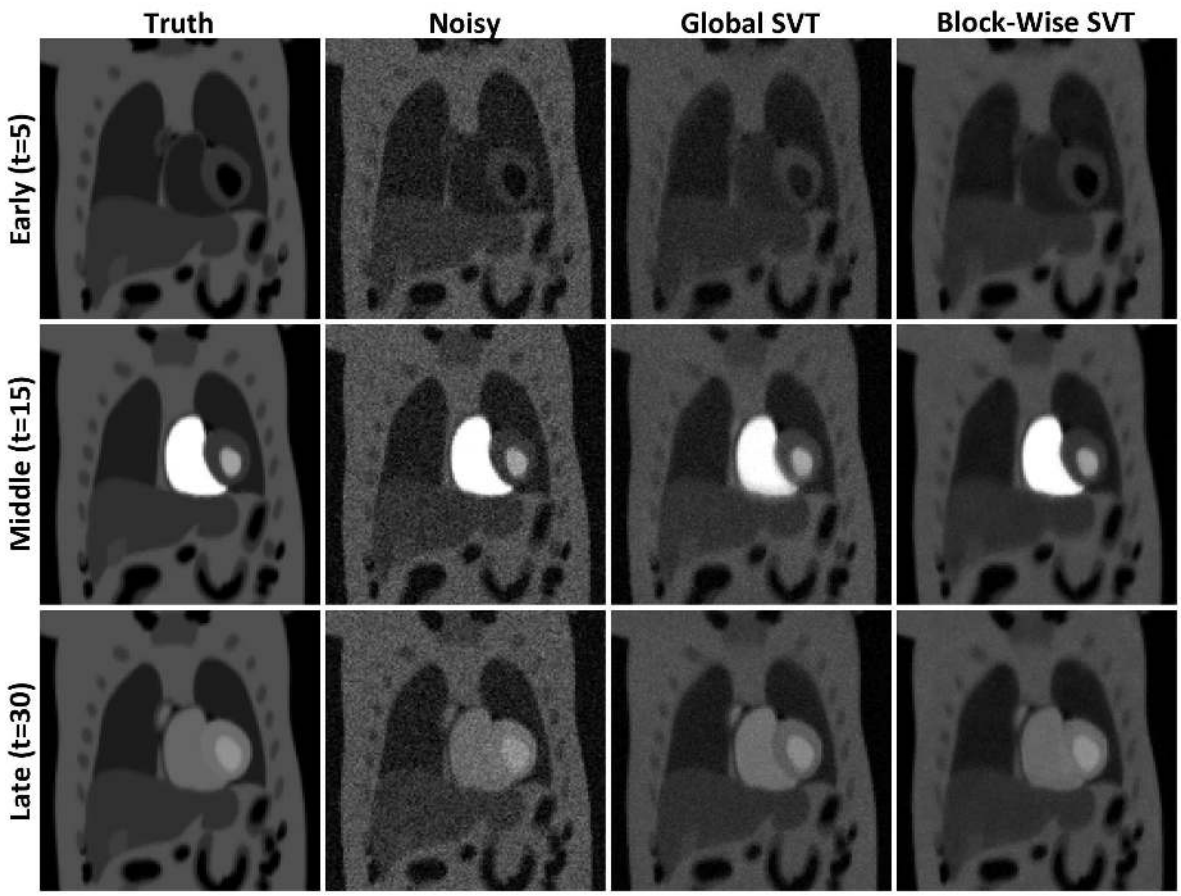

Figure 3: Comparison of MSE-optimized global and block-wise SVT for denoising the complex-valued numerical PINCAT phantom. All images are identically windowed and leveled.

and the singular value decomposition $\boldsymbol{R}_b\boldsymbol{Y} = \boldsymbol{U}_b\boldsymbol{\Sigma}_b\boldsymbol{V}_b^*$. An unbiased estimator of (2.4) is

$$\mathrm{SURE}(\mathrm{BSVT}_\lambda)(\boldsymbol{Y}) = -\beta mn\tau^2 + c^{-2}\left\|\sum_{b\in\Omega}\boldsymbol{R}_b^*\boldsymbol{U}_b\mathcal{H}_\lambda(\boldsymbol{\Sigma}_b)\boldsymbol{V}_b^*\right\|_F^2 + 2\tau^2c^{-1}\sum_{b\in\Omega}\operatorname{div}\mathrm{SVT}_\lambda(\boldsymbol{R}_b\boldsymbol{Y}), \quad (2.5)$$

where $\mathcal{H}(\boldsymbol{\Sigma}_b)_{ij} = \min(\lambda, \boldsymbol{\Sigma}_{ij}) = \mathbb{I}(i=j)\min(\lambda,\sigma_i)$ and $\sigma_i = (\Sigma_b)_{ii}$. For $\boldsymbol{Y} \in \mathbb{R}^{n^2\times t}$, $\beta = 1$; otherwise, for $\boldsymbol{Y} \in \mathbb{C}^{n^2\times t}$, $\beta = 2$.

We now present three MRI examples, one on simulated data and two on real clinical data, and demonstrate the utility of the developed unbiased risk estimators for automatic parameterization of SVT-based denoising.

Example 1: PINCAT Numerical Phantom

Initial evaluation of SURE for SVT was performed on the physiologically-improved NCAT (PINCAT) numerical phantom [31], which simulates a first-pass myocardial perfusion real-time MRI series. In particularly, the free-breathing model ($n = 128$, $t = 50$) available in the kt-SLR software

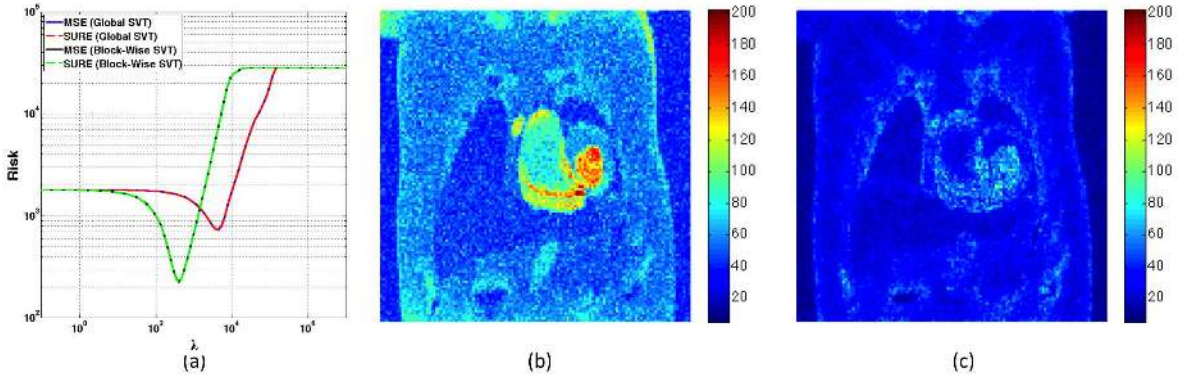

Figure 4: (a) Plots of MSE and SURE for global and block-wise SVT as a function of threshold value, λ; the minimizing parameters were used to generate the images in Fig. 3. (b) and (c) show the worst-case error through time of the MSE/SURE-optimized global and block-wise SVT results, respectively.

package [19] was adapted to include a spatially-smooth and temporally-varying phase field such that the target signal was complex valued. Complex iid Gaussian noise ($\tau = 30$) was then added to the image data. Both standard, or global, and block-wise SVT were each executed at 101 values equispaced over $[10^{-1}, 10^{7}]$. Block-wise SVT was performed with $k = 7$ and Ω comprising one block for each image pixel, under periodic boundary conditions, such that $|\Omega| = n^2$.

Figure 3 shows early, middle, and late time frames from the PINCAT series for the noise-free (truth), noisy, and SVT-denoising results. The threshold values used to generate the SVT results were selected as the MSE/SURE-minimizers in Figure 4a. Also observe in Figure 4a that SURE provides a precise estimate of MSE for both the global and block-wise SVT models. The high accuracy exhibited in this case can be attributed to the high dimensionality of the MRI series data. Note that both global and block-wise SVT yield strong noise reduction generally preserve both morphology and contrast. However, block-wise SVT simultaneously demonstrates a greater degree of noise removal and fidelity to ground truth. In particular, note the relative contrast of the various components of the heart in late frame results. The first observation is corroborated by Figure 4a, which shows that block-wise SVT is able to achieve a lower MSE than global SVT. The second observation is corroborated by Figures 4b-c, which show the worst-case absolute error (compared to ground truth) through time for the two SVT setups. Clearly, global SVT exhibits higher residual error than block-wise SVT, particularly in areas of high motion near the myocardium. The difference between these results can be attributed to the matrix anisotropy problem discussed earlier in this section. Thus, SURE can also be used to automatically-determine the block-size setting as well as the threshold value.

Example 2: Cine Cardiac Imaging

In the second experiment, a bSSFP long-axis cine sequence ($n = 192$, $t = 19$) acquired at 1.5 T using an phased-array cardiac receiver coil ($l = 8$ channels) was denoised via SVT. The Casorati matrices for individual data channels, which have undergone pre-whitening, are stacked to form a single $ln^2 \times t$ matrix. Assuming that the spatial sensitivity of the receiver channels does not vary substantially through time, the rank of this composite matrix will be equal to that for any

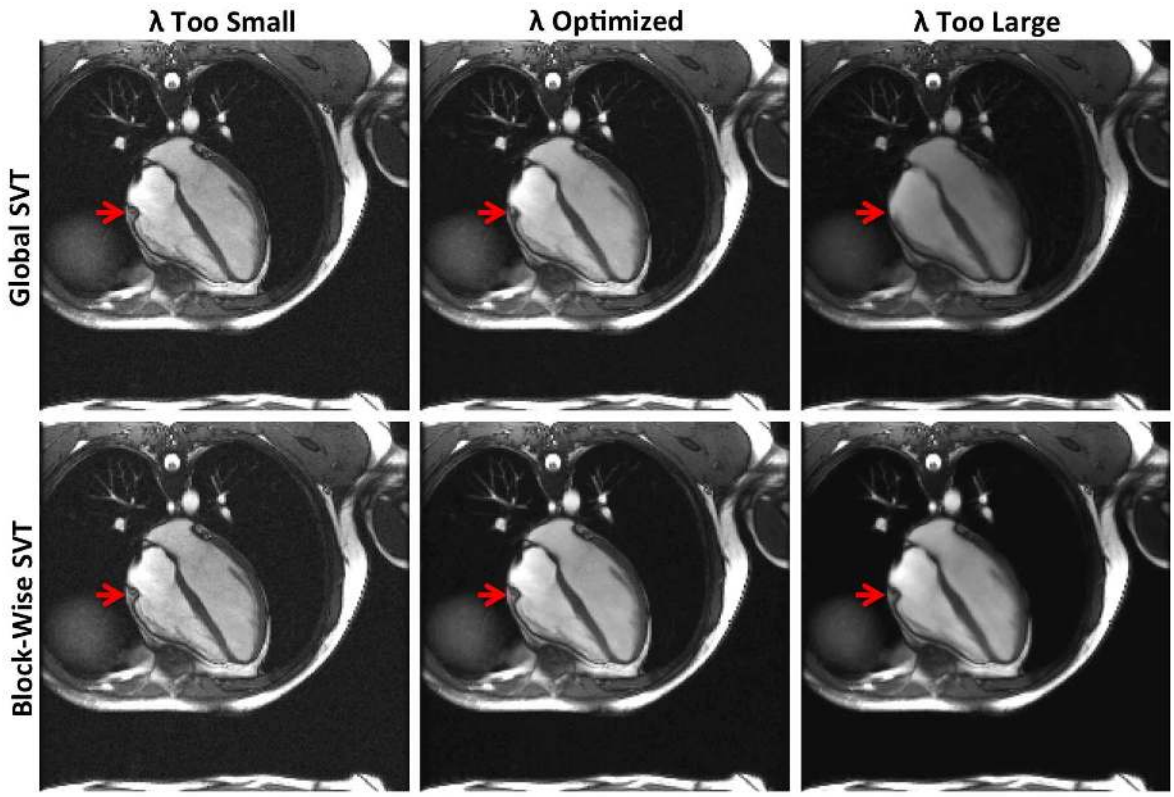

Figure 5: Global and block-wise SVT denoising results for a bSSFP long-axis cine cardiac MRI series. The effect of selecting too high or low a threshold value, as compared to the MSE/SURE optimal value is demonstrated. The red arrow identifies the tricuspid valve of the heart, which exhibit a large degree of motion through the sequence. All images are identically windowed and leveled.

individual channel. For visualization purposes, multi-channel denoising results were compressed by calculating the root-sum-of-squares image across the channel dimension. Background noise was determined using a manually-drawn region-of-interest (ROI) to have $\tau = 0.67$. In this example, block-wise SVT was performed with $k = 5$ and used the same block set as in Example 1. The parameter sweep was executed for 101 values equispaced over $[10^{-3}, 10^{5}]$. In this example, the ground truth is unknown, and thus the MSE cannot be computed to be compared against SURE. However, inspection of Figure 6a reveals that the same qualitative behavior seen for SURE in the numerical phantom example is observed here as well.

Figure 5 demonstrates the effect of sub-optimally selecting the threshold value for both global and block-wise SVT. In particular, over- and under-estimates, along with the SURE-optimal values are applied within SVT. For both SVT strategies, under-estimation of λ fails to remove noise as expected. Conversely, over-estimation of λ leads to spatio-temporal blurring by both strategies albeit in different manners. In global SVT, temporal blurring occurs near areas of high motion like the tricuspid valve of the heart (indicated via the red arrow). However, less active areas like the pulmonary branching vessels, seen just above the heart, are undistorted. Some background noise also remains visible. In the block-wise SVT result, these characteristics are essentially reversed—

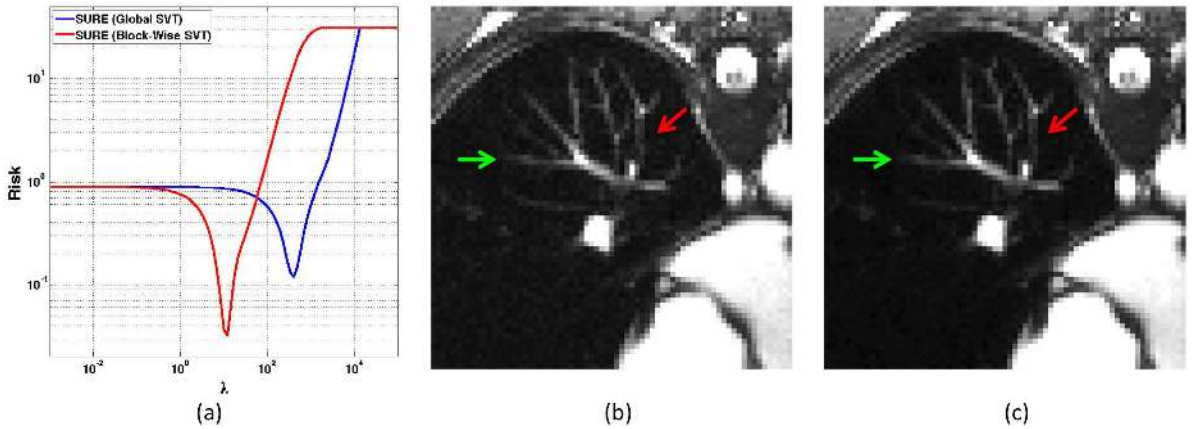

Figure 6: (a) SURE for global and block-wise SVT as a function of threshold value, λ; the minimizing parameters were used to generate the images in Fig. 5. (b) and (c) show enlargements of the pulmonary branching vessels for the SURE-optimized global and block-wise SVT results in Fig. 5, respectively. Red and green arrows highlight improved vessel conspicuity. Both images are windowed and leveled identically.

minimal temporal blurring is induced but there is noticeable spatial blurring and loss of low-contrast static features. Noise is, however, strongly reduced. A compromise between these extremes is made by the SURE-optimized results. As suggested by Figure 6a, block-wise SVT offers stronger noise suppression (see Figures 6b-c) without inducing spatial or temporal blur, the latter which is seen even in the optimized global SVT result. This example highlights the sensitivity of SVT-denoising performance on parameter selection, and the importance of having a disciplined framework for choosing threshold values. Also of note is that, even after empirical pre-whitening, multi-channel MRI noise may not be *exactly* iid Gaussian. Nonetheless, SURE allows production of extremely reliable results.

Example 3: First-Pass Myocardial Perfusion

The third denoising experiment was performed on the single-channel, first-pass myocardial perfusion real-time MRI sequence ($n_x = 190$, $n_y = 90$, $t = 70$) that is also provided with the kt-SLR software package [19]. To simulate a low-field acquisition (1.0 T), such as with an open or interventional MRI system, complex Gaussian noise was added to the data originally acquired at 3.0 T using a GRE sequence. Following this addition, the noise level was estimated at $\tau = 0.105$. Since the duration of this exam exceeded feasible breath-hold time, there is substantial respiratory motion at the end of the series. For this example, only block-wise SVT ($k = 6$) denoising was performed, and executed at 101 threshold values equispaced over $[10^{-3}, 10^{5}]$. As in Example 2, Ω comprises one block for each image pixel, and only SURE can be computed due to the absence of ground truth.

Mirroring Figure 5, Figure 7 demonstrates the effect of threshold selection on different image frames from the perfusion series. In particular, images showing the transition of contrast from the right to left vertical, as well as later-stage onset of myocardial blush are shown. Of the 101 tested threshold values, 5 threshold settings corresponding to 2 over-estimation, 2 under-estimation, and the SURE-optimal value are depicted (see Figure 8). Under all display settings, temporal contrast dynamics are well preserved; however, threshold setting has a marked effect on noise level, spatial

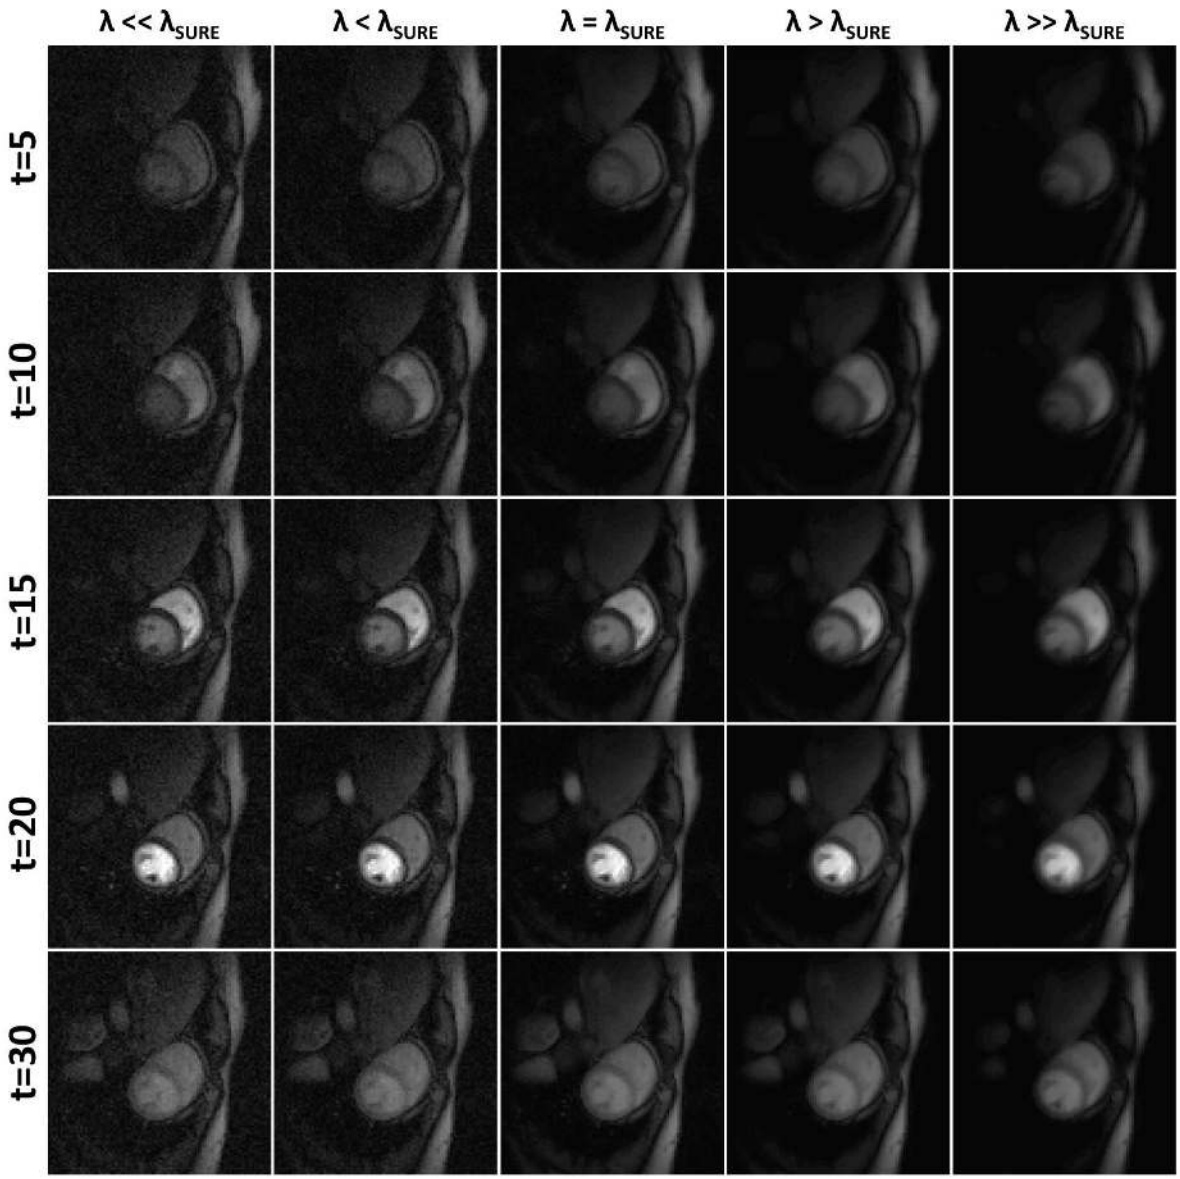

Figure 7: Block-wise SVT denoising results for GRE short-axis first-pass myocardial perfusion sequence. The effect of selecting too high or low a threshold value, as compared to the MSE/SURE optimal value, is demonstrated. All images are identically windowed and leveled.

resolution, and contrast. As before, employing SVT with an under-estimated threshold fails to remove noise, as evident in the first two columns of Figure 7. Conversely, employing SVT with an over-estimated threshold will remove noise but also induce both spatial blurring and contrast loss. This is particularly evident in the 4th image row ($t = 20$). At high threshold values, the contrast of pulmonary vasculature is diminished, and there is visible blurring of the papillary muscles, which appear as dark spots in the contrast-enhanced left ventricle. Finally, the SVT result obtained with the SURE-optimal threshold represents an ideal balance between these extremes, offering strong

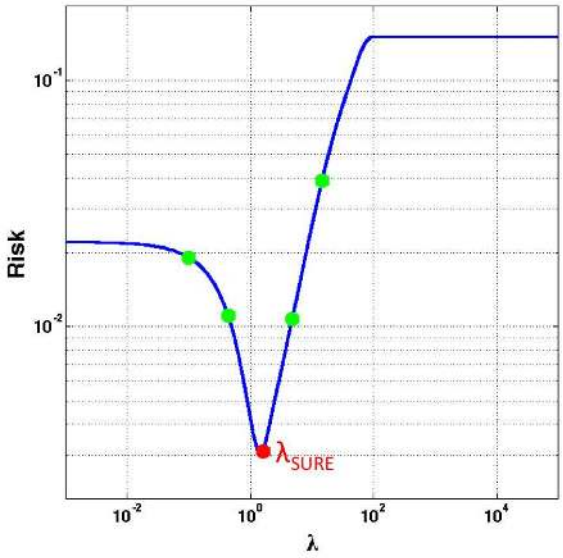

Figure 8: SURE for block-wise SVT as a function of threshold value, λ; the red dot indicates the optimal threshold value used in Fig. 7, and green dots correspond to the employed under- and over-estimated values.

noise reduction without degradation of important anatomical features. As suggested by Figure 8, there may only exist a narrow parameter window in which effective denoising can be achieved, and the presented unbiased risk estimators can greatly aid in the identification of these optimal settings.

2.2 Extensions and generalizations

The three examples in the previous subsection demonstrate the utility of the unbiased risk estimation for SVT-based denoising of cardiac MRI series data. Of course, this methodology can also be applied to any other dynamic or parametric MRI series, including those discussed earlier such as functional MRI. Although exhaustive search over a wide range of threshold values was performed for the sake of exposition, the observed unimodality of the risk functional suggests that practical automated parameter selection for SVT denoising could be accelerated using a bisection strategy such as golden section search.

In addition to optimizing over a range of threshold values, unbiased risk estimation can also be used to identify the optimal block-size when using the local SVT model in (2.1). For example, one could imagine performing a sweep of λ over a prescribed range of values for a collection of block-size values, and computing the lower envelope of this set of risk measures to investigate best achievable performance as a function of block-size. In addition to simply optimizing a denoising setup, this type of information can be used to guide the development of rank-based reconstruction frameworks for undersampled dynamic and/or parallel MRI.

3 SURE Formulas for Singular Value Thresholding

In [32], Stein gave a formula for an unbiased estimate of the mean-squared error of an estimator obeying a weak differentiability assumption and mild integrability conditions. Before reviewing this,

recall that a function $g : \mathbb{R}^p \mapsto \mathbb{R}$ is said to be weakly differentiable with respect to the variable x_i if there exists $h : \mathbb{R}^p \mapsto \mathbb{R}$ such that for all compactly supported and infinitely differentiable functions φ,

$$\int \varphi(x)h(x)\,dx = -\int \frac{\partial\varphi(x)}{\partial x_i} g(x)\,dx.$$

Roughly speaking, the derivatives can fail to exist over regions of Lebesgue measure zero.

Proposition 3.1 ([32, 14])**.** *Suppose that $Y_{ij} \overset{iid}{\sim} \mathcal{N}(X_{ij}, 1)$. Consider an estimator $\hat{\boldsymbol{X}}$ of the form $\hat{\boldsymbol{X}} = \boldsymbol{Y} + g(\boldsymbol{Y})$, where $g_{ij} : \mathbb{R}^{m\times n} \mapsto \mathbb{R}$ is weakly differentiable with respect to Y_{ij} and*

$$\mathbb{E}\left\{|Y_{ij}g_{ij}(\boldsymbol{Y})| + \left|\frac{\partial}{\partial Y_{ij}} g_{ij}(\boldsymbol{Y})\right|\right\} < \infty,$$

for $(i, j) \in \mathcal{I} := \{1, \dots, m\} \times \{1, \dots n\}$. Then

$$\mathbb{E}\,\|\hat{\boldsymbol{X}} - \boldsymbol{X}\|_F^2 = \mathbb{E}\big\{mn + 2\mathrm{div}\,(g(\boldsymbol{Y})) + \|g(\boldsymbol{Y})\|_F^2\big\}. \tag{3.1}$$

The rest of this section establishes that SVT obeys these assumptions whereas the deduction of a closed-form expression for its divergence is deferred to Section 4. To achieve this, we record a simple fact that will allow us to study both the real- and complex-valued cases within the same framework. We identify below $\mathbb{R}^{m\times n}$ with $\mathbb{R}^{mn}$.

Lemma 3.2. *Let $g : \mathbb{R}^{m\times n} \mapsto \mathbb{R}^{m\times n}$ with components $g_{ij} : \mathbb{R}^{m\times n} \mapsto \mathbb{R}$. Assume $g(\boldsymbol{0}) = \boldsymbol{0}$. If g is Lipschitz with constant $L > 0$, i.e.,*

$$\forall\, \boldsymbol{X}, \boldsymbol{Y} : \quad \|g(\boldsymbol{X}) - g(\boldsymbol{Y})\| \le L\|\boldsymbol{X} - \boldsymbol{Y}\|,$$

for some norm in $\mathbb{R}^{m\times n}$, then each g_{ij} is weakly differentiable Lebesgue-a.e. in $\mathbb{R}^{m\times n}$. Furthermore, if $\boldsymbol{Y} \in \mathbb{R}^{m\times n}$ is distributed as in (1.1)*, then for all pairs (i, j),*

$$\mathbb{E}\,\{|Y_{ij}g_{ij}(\boldsymbol{Y})|\} < \infty \quad \text{and} \quad \mathbb{E}\left\{\left|\frac{\partial}{\partial Y_{ij}} g_{ij}(\boldsymbol{Y})\right|\right\} < \infty.$$

Proof Since all norms are equivalent in finite dimensions, g is also Lipschitz with respect to the Frobenius norm. Hence,

$$\forall\, \boldsymbol{X}, \boldsymbol{Y} : \quad \|g(\boldsymbol{X}) - g(\boldsymbol{Y})\|_F \le L\|\boldsymbol{X} - \boldsymbol{Y}\|_F \quad \Longrightarrow \quad |g_{ij}(\boldsymbol{X}) - g_{ij}(\boldsymbol{Y})| \le L\|\boldsymbol{X} - \boldsymbol{Y}\|_F, \tag{3.2}$$

for some constant L. By Rademacher's theorem (see Section 3.1.2 in [12]) each component is differentiable Lebesgue-a.e. in $\mathbb{R}^{m\times n}$. Furthermore, these components are weakly differentiable (see Theorem 1 and 2, Section 6.2, in [12]), and both the derivatives and weak derivatives are Lebesgue-a.e. equal.

Regarding the integrability, since $g(\boldsymbol{0}) = \boldsymbol{0}$, (3.2) yields $\|g(\boldsymbol{Y})\|_F^2 \le L^2\|\boldsymbol{Y}\|_F^2$ and we deduce

$$\mathbb{E}\,\|g(\boldsymbol{Y})\|_F^2 \le L^2\,\mathbb{E}\,\|\boldsymbol{Y}\|_F^2 = L^2(mn + \|\boldsymbol{X}_0\|_F^2) < \infty.$$

Furthermore, the Cauchy-Schwarz inequality gives

$$\mathbb{E}\{|Y_{ij}g_{ij}(\boldsymbol{Y})|\} \le \mathbb{E}\{Y_{ij}^2\}^{1/2}\,\mathbb{E}\{g_{ij}(\boldsymbol{Y})^2\}^{1/2} \le L\left[(1 + X_{0,ij}^2)(mn + \|\boldsymbol{X}_0\|_F^2)\right]^{1/2} < \infty.$$

Finally, (3.2) asserts that the derivatives of g_{ij}—whenever they exist—are bounded by L. Hence,

$$\mathbb{E}\left\{\left|\frac{\partial}{\partial Y_{ij}} g_{ij}(\boldsymbol{Y})\right|\right\} \leq \mathbb{E}\left\{\left|\frac{\partial}{\partial Y_{ij}} g_{ij}(\boldsymbol{Y})\right|^2\right\}^{1/2} \leq L,$$

which concludes the proof. ∎

For pedagogical reasons, we work with singular value thresholding first, before discussing more general spectral estimators in Section 4. Writing $g(\boldsymbol{Y}) = \mathrm{SVT}_\lambda(\boldsymbol{Y}) - \boldsymbol{Y}$ then, the conditions become

$$\mathbb{E}\left\{|Y_{ij} g_{ij}(\boldsymbol{Y})| + \left|\frac{\partial g_{ij}(\boldsymbol{Y})}{\partial Y_{ij}}\right|\right\} \leq \mathbb{E}\left\{|Y_{ij}\mathrm{SVT}_\lambda(\boldsymbol{Y})_{ij}| + \left|\frac{\partial}{\partial Y_{ij}}\mathrm{SVT}_\lambda(\boldsymbol{Y})_{ij}\right|\right\} + \mathbb{E}\{Y_{ij}^2\} + 1,$$

and we only need to study the weak differentiability and integrability of SVT.

3.1 The real case

Lemma 3.3. *The mapping* SVT_λ, $\lambda > 0$, *obeys the assumptions of Proposition 3.1.*

Proof By definition, SVT_λ is the proximity function of the nuclear norm. Since the nuclear norm is a convex, proper and lower semi-continuous function, SVT_λ is Lipschitz and non-expansive (see Chapter 6, Section 30 in [30]). Applying Lemma 3.2 proves the claim. ∎

The mapping SVT_λ may not be differentiable everywhere. Fortunately, since we are interested in weak differentiability, we can discard sets of Lebesgue measure zero. For instance, model (1.1) guarantees that $\boldsymbol{Y}$ is simple and has full-rank with probability one. Further, since SVT_λ acts by applying soft-thresholding to each singular value, it seems natural that matrices with at least one singular value equal to λ will be a cause of concern. Consider then

$$\mathcal{F} := \{\boldsymbol{X} \in \mathbb{R}^{m\times n} : \boldsymbol{X} \text{ is simple, has full rank and no singular value exactly equal to } \lambda\}.$$

It is clear that the complement of $\mathcal{F}$ has Lebesgue measure zero. Consequently, we can restrict ourselves to the open set $\mathcal{F}$ to look for a divergence. The result below—a corollary of the general Theorem 4.3—establishes that SVT is differentiable over $\mathcal{F}$, and thus we have a divergence in the usual sense.

Corollary 3.4. *The mapping* SVT_λ *is differentiable over* $\mathcal{F}$ *and we have*

$$\mathrm{div}\left(\mathrm{SVT}_\lambda(\boldsymbol{X})\right) = \sum_{i=1}^{\min(m,n)} \left[\mathbb{I}\{\sigma_i > \lambda\} + |m-n|\left(1 - \frac{\lambda}{\sigma_i}\right)_+\right] + 2\sum_{i\neq j,\, i,j=1}^{\min(m,n)} \frac{\sigma_i(\sigma_i - \lambda)_+}{\sigma_i^2 - \sigma_j^2}. \tag{3.3}$$

Since SVT_λ is of the form (1.11) with

$$f_i(\sigma) = (\sigma - \lambda)_+, \quad f_i'(\sigma) = \mathbb{I}\{\sigma > \lambda\}, \tag{3.4}$$

for $i = 1, \ldots, \min(m,n)$ and $\sigma \neq \lambda$, we see that it is differentiable at the singular value matrix of any element of $\mathcal{F}$. Because the elements of $\mathcal{F}$ are also simple and full-rank, we conclude SVT_λ is differentiable over $\mathcal{F}$ by Lemma 4.1. Therefore, the proof of (3.3) is a special case of the general Theorem 4.3 from Section 4.

3.2 The complex case

In the complex case, we need to analyze the conditions for SURE with respect to model (1.9), which can be expressed in vector form as

$$\begin{bmatrix}\mathrm{Re}(\boldsymbol{Y})\\ \mathrm{Im}(\boldsymbol{Y})\end{bmatrix} = \begin{bmatrix}\mathrm{Re}(\boldsymbol{X}_0)\\ \mathrm{Im}(\boldsymbol{X}_0)\end{bmatrix} + \begin{bmatrix}\mathrm{Re}(\boldsymbol{W})\\ \mathrm{Im}(\boldsymbol{W})\end{bmatrix}. \tag{3.5}$$

We thus need to study the existence of the weak partial derivatives

$$\frac{\partial}{\partial \mathrm{Re}(Y_{ij})}\mathrm{Re}(\mathrm{SVT}_\lambda)_{ij} \qquad \text{and} \qquad \frac{\partial}{\partial \mathrm{Im}(Y_{ij})}\mathrm{Im}(\mathrm{SVT}_\lambda)_{ij},$$

and whether they satisfy our integrability conditions, as these are the only ones that play a role in the divergence.

Lemma 3.5. *The mappings* SVT_λ, $\mathrm{Re}(\mathrm{SVT}_\lambda)$ *and* $\mathrm{Im}(\mathrm{SVT}_\lambda)$ *obey the assumptions of Proposition 3.1.*

Proof As in the real case, SVT_λ is the prox-function of the nuclear norm and is thus Lipschitz and non-expansive (see Proposition 2.27 in [1]). Hence, the real and imaginary parts are Lipschitz. Applying Lemma 3.2 to both the real and imaginary components proves the claim. ■

For the divergence, we write

$$\mathrm{div}\left(\mathrm{SVT}_\lambda(\boldsymbol{X})\right) = \mathrm{div}_{\mathrm{Re}(\boldsymbol{X})}\left(\mathrm{Re}(\mathrm{SVT}_\lambda)(\boldsymbol{X})\right) + \mathrm{div}_{\mathrm{Im}(\boldsymbol{X})}\left(\mathrm{Im}(\mathrm{SVT}_\lambda)(\boldsymbol{X})\right). \tag{3.6}$$

As in the real case, we can discard sets of Lebesgue measure zero. Consider then the analogue of $\mathcal{F}$, whose complement has Lebesgue measure zero.

Corollary 3.6. *The mapping* SVT_λ *is differentiable over* $\mathcal{F}$*, and we have*

$$\mathrm{div}\left(\mathrm{SVT}_\lambda(\boldsymbol{X})\right) = \sum_{i=1}^{\min(m,n)}\left[\mathbb{I}\{\sigma_i > \lambda\} + (2|m-n|+1)\left(1-\frac{\lambda}{\sigma_i}\right)_+\right] + 4\sum_{i\neq j,\, i,j=1}^{\min(m,n)}\frac{\sigma_i(\sigma_i-\lambda)_+}{\sigma_i^2-\sigma_j^2}. \tag{3.7}$$

The argument regarding the differentiability of SVT_λ applies here as well. The proof of the closed-form expression for the divergence is a special case of the general Theorem 4.5 from Section 4 with f_i as in (3.4).

4 Differentiability of Spectral Functions: General SURE Formulas

This section introduces some results concerning the differentiability of spectral functions.[2] To simplify the exposition, we work with $m \times n$ matrices $\boldsymbol{X}$ with $m \geq n$, as all our arguments apply with minor modifications to $\boldsymbol{X}^*$. Below, $\{\boldsymbol{E}^{ij}\}_{(i,j)\in\mathcal{I}}$ is the canonical basis of $\mathbb{R}^{n\times m}$, $\boldsymbol{I}_{n\times n}$ is the n-dimensional identity matrix and we set

$$\boldsymbol{I}_{m\times n} = \begin{bmatrix}\boldsymbol{I}_{n\times n}\\ \boldsymbol{0}\end{bmatrix}.$$

[2]Here, we set aside any statistical interpretation to focus on differentiability and closed-form expression for the divergence of spectral functions.

In the real and complex cases, we make use of the reduced and full SVD. The reduced SVD is written as $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^*$, where $\boldsymbol{U}$ is $m \times n$ with orthonormal columns and $\boldsymbol{V}$ is unitary. In the full SVD, $\boldsymbol{X} = \tilde{\boldsymbol{U}}\boldsymbol{\Sigma}\boldsymbol{V}^*$ where $\tilde{\boldsymbol{U}}$ is $m \times m$ and extends $\boldsymbol{U}$ to a unitary matrix. Finally, we recall from Section 1 that a spectral function is of the form $f(\boldsymbol{X}) = \boldsymbol{U} f(\boldsymbol{\Sigma})\boldsymbol{V}^*$ with

$$f(\boldsymbol{\Sigma}) = \operatorname{diag}(f_1(\sigma_1), \ldots, f_n(\sigma_n)), \tag{4.1}$$

where we assume that each $f_i : \mathbb{R}_+ \mapsto \mathbb{R}_+$ is differentiable. We note that we can relax the differentiability assumption and only require differentiability in a neighborhood of the spectrum of the matrix under study.

To find a closed-form for the divergence of f, we first need to determine whether f is differentiable. Using notation from differential calculus (see [10]), recall that f is differentiable at $\boldsymbol{X}$ if there exists a linear mapping $\mathrm{d}f_{\boldsymbol{X}}$, called the differential of f at $\boldsymbol{X}$, such that for all $\boldsymbol{\Delta}$,

$$\lim_{\boldsymbol{\Delta} \to \boldsymbol{0}} \frac{\|f(\boldsymbol{X} + \boldsymbol{\Delta}) - f(\boldsymbol{X}) - \mathrm{d}f_{\boldsymbol{X}}[\boldsymbol{\Delta}]\|_F}{\|\boldsymbol{\Delta}\|_F} = 0.$$

Differentiation obeys the standard rules of calculus and we have

$$\mathrm{d}f[\boldsymbol{\Delta}] = \mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}]\, f(\boldsymbol{\Sigma})\, \boldsymbol{V}^* + \boldsymbol{U}\, \mathrm{d}(f \circ \boldsymbol{\Sigma})[\boldsymbol{\Delta}]\, \boldsymbol{V}^* + \boldsymbol{U}\, f(\boldsymbol{\Sigma})\, \mathrm{d}\boldsymbol{V}[\boldsymbol{\Delta}]^*. \tag{4.2}$$

where, to lighten the notation above, we have not indicated the point $\boldsymbol{X}$ at which the differential is calculated (we shall use subscripts whenever necessary). By the chain rule, $\mathrm{d}(f \circ \boldsymbol{\Sigma})_{\boldsymbol{X}}[\boldsymbol{\Delta}] = \mathrm{d}f_{\boldsymbol{\Sigma}}[\mathrm{d}\boldsymbol{\Sigma}_{\boldsymbol{X}}[\boldsymbol{\Delta}]]$ so that the differentiability of f depends upon the existence of differentials for $\boldsymbol{U}$, $\boldsymbol{V}$ and $\boldsymbol{\Sigma}$ at $\boldsymbol{X}$ together with that of f at $\boldsymbol{\Sigma}$.

Standard analysis arguments immediately establish that the function f is differentiable at $\boldsymbol{\Sigma}$ whenever $\boldsymbol{\Sigma}$ is simple, as in this case there is no ambiguity as to which function f_i is being applied to the singular values. Furthermore, the differentiability of the SVD of a simple matrix with full-rank is a consequence of Theorem 1 and Theorem 2 in [23] applied to $\boldsymbol{X}^*\boldsymbol{X}$ and $\boldsymbol{X}\boldsymbol{X}^*$ (see also [16] and [33]). We summarize these arguments in the following lemma.

Lemma 4.1. *Let $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^*$ be a simple and full-rank matrix and f be a spectral function.*

1. *The factors of the SVD are differentiable in a neighborhood of $\boldsymbol{X}$.*
2. *If f is differentiable at $\boldsymbol{\Sigma}$, then it is differentiable at $\boldsymbol{X}$.*

We now focus on determining the divergence of f in closed-form. Since there are differences between the real- and complex-valued cases, we treat them separately.

4.1 The real case

Lemma 4.1 asserts that the SVD is differentiable at any simple matrix $\boldsymbol{X}$ with full-rank. The following result shows that we can determine the differentials of each factor in closed-form.

Lemma 4.2. *Let $\boldsymbol{X}$ be simple and full-rank. Then the differentials of $\boldsymbol{U}$ and $\boldsymbol{V}$ are given by*

$$\mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}] = \tilde{\boldsymbol{U}}\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}], \qquad \mathrm{d}\boldsymbol{V}[\boldsymbol{\Delta}] = \boldsymbol{V}\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}]^*,$$

where $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}$ and $\boldsymbol{\Omega}_{\boldsymbol{V}}$ are given in closed-form (4.8) *($\mathrm{d}\boldsymbol{U}$ is independent of the choice of $\tilde{\boldsymbol{U}}$). Finally, $\mathrm{d}\boldsymbol{\Sigma}$ is given by* (4.7)*.*

Proof We follow the same method as in [10], which also appears in [27]. This method is also that used in [8] as mentioned earlier. Let $\boldsymbol{X} \in \mathbb{R}^{m\times n}$ $(m \geq n)$. For any matrix $\boldsymbol{\Delta} \in \mathbb{R}^{m\times n}$, we have

$$\mathrm{d}\boldsymbol{X}[\boldsymbol{\Delta}] = \mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}]\boldsymbol{\Sigma}\boldsymbol{V}^* + \boldsymbol{U}\mathrm{d}\boldsymbol{\Sigma}[\boldsymbol{\Delta}]\boldsymbol{V}^* + \boldsymbol{U}\boldsymbol{\Sigma}\mathrm{d}\boldsymbol{V}[\boldsymbol{\Delta}]^*. \tag{4.3}$$

Since $\tilde{\boldsymbol{U}}$ and $\boldsymbol{V}$ are orthogonal, it follows that

$$\tilde{\boldsymbol{U}}^*\mathrm{d}\boldsymbol{X}[\boldsymbol{\Delta}]\boldsymbol{V} = \tilde{\boldsymbol{U}}^*\mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}]\boldsymbol{\Sigma} + \tilde{\boldsymbol{U}}^*\boldsymbol{U}\mathrm{d}\boldsymbol{\Sigma}[\boldsymbol{\Delta}] + \tilde{\boldsymbol{U}}^*\boldsymbol{U}\boldsymbol{\Sigma}\mathrm{d}\boldsymbol{V}[\boldsymbol{\Delta}]^*\boldsymbol{V}.$$

Introduce the matrices

$$\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}] = \tilde{\boldsymbol{U}}^*\mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}] \qquad \text{and} \qquad \boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}] = \mathrm{d}\boldsymbol{V}[\boldsymbol{\Delta}]^*\boldsymbol{V},$$

so that

$$\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V} = \boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}]\boldsymbol{\Sigma} + \boldsymbol{I}_{m\times n}\mathrm{d}\boldsymbol{\Sigma}[\boldsymbol{\Delta}] + \boldsymbol{I}_{m\times n}\boldsymbol{\Sigma}\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}], \tag{4.4}$$

where we used the fact that $\mathrm{d}\boldsymbol{X}$ is the identity, whence $\mathrm{d}\boldsymbol{X}[\boldsymbol{\Delta}] = \boldsymbol{\Delta}$. Moreover,

$$\boldsymbol{0} = \mathrm{d}\boldsymbol{I}_{n\times n}(\boldsymbol{\Delta}) = \boldsymbol{V}^*\mathrm{d}\boldsymbol{V}[\boldsymbol{\Delta}] + \mathrm{d}\boldsymbol{V}[\boldsymbol{\Delta}]^*\boldsymbol{V} \quad \Longrightarrow \quad \boldsymbol{V}^*\mathrm{d}\boldsymbol{V}[\boldsymbol{\Delta}] = -\mathrm{d}\boldsymbol{V}[\boldsymbol{\Delta}]^*\boldsymbol{V},$$

which says that $\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}]$ is anti-symmetric. Similarly,

$$\boldsymbol{0} = \mathrm{d}\boldsymbol{I}_{n\times n}(\boldsymbol{\Delta}) = \boldsymbol{U}^*\mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}] + \mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}]^*\boldsymbol{U} \quad \Longrightarrow \quad \boldsymbol{U}^*\mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}] = -\mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}]^*\boldsymbol{U},$$

and thus the upper $n \times n$ block of $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}]$ is also anti-symmetric. This fact and the skew-symmetry of $\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}]$ yield the relations

$$\begin{aligned}
\Omega_{\tilde{\boldsymbol{U}},ii}[\boldsymbol{\Delta}] &= 0, \qquad i = 1,\ldots,n, \\
\Omega_{\boldsymbol{V},ii}[\boldsymbol{\Delta}] &= 0, \qquad i = 1,\ldots,n, \\
(\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ij} &= \Omega_{\tilde{\boldsymbol{U}},ij}[\boldsymbol{\Delta}]\sigma_j + \sigma_i\Omega_{\boldsymbol{V},ij}[\boldsymbol{\Delta}], \qquad i \neq j, \quad i,j = 1,\ldots,n, \qquad (4.5) \\
-(\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ji} &= \Omega_{\tilde{\boldsymbol{U}},ij}[\boldsymbol{\Delta}]\sigma_i + \sigma_j\Omega_{\boldsymbol{V},ij}[\boldsymbol{\Delta}], \qquad i \neq j, \quad i,j = 1,\ldots,n, \qquad (4.6)
\end{aligned}$$

and

$$\mathrm{d}\sigma_i[\boldsymbol{\Delta}] = (\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ii}, \qquad i = 1,\ldots,n. \tag{4.7}$$

In particular, (4.5) and (4.6) can be summarized in the system of linear equations

$$\begin{bmatrix} \sigma_j & \sigma_i \\ \sigma_i & \sigma_j \end{bmatrix} \begin{bmatrix} \Omega_{\tilde{\boldsymbol{U}},ij}[\boldsymbol{\Delta}] \\ \Omega_{\boldsymbol{V},ij}[\boldsymbol{\Delta}] \end{bmatrix} = \begin{bmatrix} (\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ij} \\ -(\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ji} \end{bmatrix}, \qquad i \neq j, \quad i,j = 1,\ldots,n.$$

Since $\boldsymbol{X}$ is simple, the coefficient matrix is invertible, and thus

$$\begin{bmatrix} \Omega_{\tilde{\boldsymbol{U}},ij}[\boldsymbol{\Delta}] \\ \Omega_{\boldsymbol{V},ij}[\boldsymbol{\Delta}] \end{bmatrix} = -\frac{1}{\sigma_i^2 - \sigma_j^2} \begin{bmatrix} \sigma_j & \sigma_i \\ -\sigma_i & -\sigma_j \end{bmatrix} \begin{bmatrix} (\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ij} \\ (\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ji} \end{bmatrix}, \qquad i \neq j, \quad i,j = 1,\ldots,n. \tag{4.8}$$

These relations determine both $\mathrm{d}\boldsymbol{\Sigma}[\boldsymbol{\Delta}]$, $\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}]$ and the upper $n \times n$ block of $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}]$ completely. The lower $(m-n) \times n$ block of $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}]$ is determined from (4.4), as

$$(\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ij} = \Omega_{\tilde{\boldsymbol{U}},ij}[\boldsymbol{\Delta}]\sigma_j, \qquad i = n+1,\ldots,m, \quad j = 1,\ldots,n. \tag{4.9}$$

Since $\boldsymbol{X}$ has full-rank, this completely determines $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}]$. We conclude that (i) $\mathrm{d}\boldsymbol{\Sigma}[\boldsymbol{\Delta}]$, $\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}]$ and $\boldsymbol{\Omega}_{\boldsymbol{U}}[\boldsymbol{\Delta}]$ are uniquely characterized by the SVD of $\boldsymbol{X}$, and (ii) they act linearly on $\boldsymbol{\Delta}$. Furthermore, the differentials obey (4.3). We now prove that $\mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}]$ is independent of the choice of $\tilde{\boldsymbol{U}}$. Decompose $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}]$ and $\tilde{\boldsymbol{U}}$ as

$$\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}] = \begin{bmatrix} \boldsymbol{\Omega}_{\boldsymbol{U}}[\boldsymbol{\Delta}] \\ \boldsymbol{\Omega}_{\boldsymbol{Q}}[\boldsymbol{\Delta}] \end{bmatrix} \qquad \text{and} \qquad \tilde{\boldsymbol{U}} = \begin{bmatrix} \boldsymbol{U} & \boldsymbol{Q} \end{bmatrix},$$

where $\boldsymbol{\Omega}_{\boldsymbol{U}}[\boldsymbol{\Delta}]$ denotes the upper $n \times n$ block of $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}]$ (which depends only on $\boldsymbol{U}$ by (4.8)) and $\boldsymbol{\Omega}_{\boldsymbol{Q}}[\boldsymbol{\Delta}]$ denotes the lower $(m-n) \times n$ block of $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}]$ (which by (4.9) equals $\boldsymbol{Q}^*\boldsymbol{\Delta}\boldsymbol{V}\boldsymbol{\Sigma}^{-1}$). Then

$$\mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}] = \tilde{\boldsymbol{U}}\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}] = \boldsymbol{U}\boldsymbol{\Omega}_{\boldsymbol{U}}[\boldsymbol{\Delta}] + \boldsymbol{Q}\boldsymbol{\Omega}_{\boldsymbol{Q}}[\boldsymbol{\Delta}] = \boldsymbol{U}\boldsymbol{\Omega}_{\boldsymbol{U}}[\boldsymbol{\Delta}] + (\boldsymbol{Q}\boldsymbol{Q}^*)\boldsymbol{\Delta}\boldsymbol{V}\boldsymbol{\Sigma}^{-1},$$

which is well defined since the orthogonal projector $\boldsymbol{Q}\boldsymbol{Q}^*$ is independent of the choice of $\boldsymbol{Q}$. ∎

Since Lemma 4.2 provides a closed-form expression for the differentials, computing the value of the divergence is now a matter of calculus.

Theorem 4.3. *Let f be a matrix-valued spectral function and suppose $\boldsymbol{X} \in \mathbb{R}^{m\times n}$ is simple and has full-rank. Then*

$$\operatorname{div}(f(\boldsymbol{X})) = \sum_{i=1}^{\min(m,n)} \left(f_i'(\sigma_i) + |m-n|\frac{f_i'(\sigma_i)}{\sigma_i} \right) + 2\sum_{i\neq j,\, i,j=1}^{\min(m,n)} \frac{\sigma_i f_i(\sigma_i)}{\sigma_i^2 - \sigma_j^2}. \tag{4.10}$$

Proof Assume $m \geq n$ without loss of generality. Let $\{\tilde{\boldsymbol{u}}_i\}_{i=1}^m$ and $\{\boldsymbol{v}_i\}_{i=1}^n$ denote the columns of $\tilde{\boldsymbol{U}}$ and $\boldsymbol{V}$, and set

$$\boldsymbol{\Delta}^{ij} = \tilde{\boldsymbol{u}}_i \boldsymbol{v}_j^*, \qquad (i,j) \in \mathcal{I}.$$

Then $\{\boldsymbol{\Delta}^{ij}\}_{(i,j)\in\mathcal{I}}$ is an orthonormal basis for $\mathbb{R}^{m\times n}$ and, furthermore, $\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}^{ij}\boldsymbol{V} = \boldsymbol{E}^{ij}$. Since $\{\boldsymbol{\Delta}^{ij}\}_{(i,j)\in\mathcal{I}}$ is an orthonormal basis,

$$\operatorname{div}(f(\boldsymbol{X})) = \sum_{(i,j)\in\mathcal{I}} \left\langle \boldsymbol{\Delta}^{ij}, \mathrm{d}f[\boldsymbol{\Delta}^{ij}] \right\rangle. \tag{4.11}$$

Note that (4.2) yields

$$\mathrm{d}f[\boldsymbol{\Delta}^{ij}] = \tilde{\boldsymbol{U}}^*(\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}^{ij}]f(\boldsymbol{\Sigma}) + \boldsymbol{I}_{m\times n}\mathrm{d}(f\circ\boldsymbol{\Sigma})[\boldsymbol{\Delta}^{ij}] + \boldsymbol{I}_{m\times n}f(\boldsymbol{\Sigma})\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}^{ij}])\boldsymbol{V}^*.$$

Next, from (4.7) we obtain

$$\mathrm{d}\sigma_k[\boldsymbol{\Delta}^{ij}] = \delta_{ki}\delta_{kj}, \qquad (i,j) \in \mathcal{I}, \quad k = 1,\ldots,n.$$

By the same arguments, from (4.8) we obtain

$$\begin{bmatrix} \Omega_{\tilde{\boldsymbol{U}},kl}[\boldsymbol{\Delta}^{ij}] \\ \Omega_{\boldsymbol{V},kl}[\boldsymbol{\Delta}^{ij}] \end{bmatrix} = -\frac{1}{\sigma_k^2 - \sigma_l^2} \begin{bmatrix} \sigma_l & \sigma_k \\ -\sigma_k & -\sigma_l \end{bmatrix} \begin{bmatrix} \delta_{ik}\delta_{jl} \\ \delta_{il}\delta_{jk} \end{bmatrix}, \qquad (i,j) \in \mathcal{I}, \quad k,l = 1,\ldots,n, \quad k \neq l,$$

and from (4.9),

$$\Omega_{\tilde{\boldsymbol{U}},kl}[\boldsymbol{\Delta}^{ij}] = \frac{\delta_{ik}\delta_{jl}}{\sigma_l}, \qquad (i,j) \in \mathcal{I}, \quad k = n+1,\ldots,m, \quad l = 1,\ldots,n.$$

Now follow (4.11) and decompose the divergence as

$$\begin{aligned}\operatorname{div}(f(\boldsymbol{X})) &= \sum_{(i,j)\in\mathcal{I}} \left\langle \boldsymbol{E}^{ij}, \boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}^{ij}] f(\boldsymbol{\Sigma}) + \boldsymbol{I}_{m\times n} \mathrm{d}(f\circ\boldsymbol{\Sigma})[\boldsymbol{\Delta}^{ij}] + \boldsymbol{I}_{m\times n} f(\boldsymbol{\Sigma}) \boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}^{ij}] \right\rangle \\ &= S^{\tilde{\boldsymbol{U}}} + S^{\boldsymbol{\Sigma}} + S^{\boldsymbol{V}}.\end{aligned}$$

We see that $S^{\boldsymbol{V}}$ is equal to

$$S^{\boldsymbol{V}} = \sum_{i\neq j,\, i,j=1}^{n} f_{ij}(\boldsymbol{\Sigma}) \Omega_{\boldsymbol{V},ij}[\boldsymbol{\Delta}^{ij}] = \sum_{i\neq j,\, i,j=1}^{n} \frac{f_i(\sigma_i)}{\sigma_i^2-\sigma_j^2}(\sigma_i \delta_{ii}\delta_{jj} + \sigma_j \delta_{ij}\delta_{ji}) = \sum_{i\neq j,\, i,j=1}^{n} \frac{\sigma_i f_i(\sigma_i)}{\sigma_i^2-\sigma_j^2},$$

whereas

$$S^{\boldsymbol{\Sigma}} = \sum_{i,j=1}^{n} \mathrm{d}f_{ij}[\mathrm{d}\Sigma_{ij}[\boldsymbol{\Delta}^{ij}]] = \sum_{i,j=1}^{n} \delta_{ij} f_i'(\sigma_i) = \sum_{i=1}^{n} f_i'(\sigma_i).$$

Finally, $S^{\tilde{\boldsymbol{U}}}$ decomposes as

$$S^{\tilde{\boldsymbol{U}}} = \sum_{i,j=1}^{n} \left\langle \boldsymbol{E}^{ij}, \boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}^{ij}] f(\boldsymbol{\Sigma}) \right\rangle + \sum_{i=n+1}^{m} \sum_{j=1}^{n} \left\langle \boldsymbol{E}^{ij}, \boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}^{ij}] f(\boldsymbol{\Sigma}) \right\rangle.$$

Using (4.8), the first term is equal to

$$\sum_{i,j=1}^{n} \left\langle \boldsymbol{E}^{ij}, \boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}^{ij}] f(\boldsymbol{\Sigma}) \right\rangle = -\sum_{i\neq j,\, i,j=1}^{n} \frac{f_j(\sigma_j)}{\sigma_i^2-\sigma_j^2}(\sigma_j \delta_{ii}\delta_{jj} + \sigma_i \delta_{ij}\delta_{ji}) = \sum_{i\neq j,\, i,j=1}^{n} \frac{\sigma_i f_i(\sigma_i)}{\sigma_i^2-\sigma_j^2},$$

while it follows from (4.9) that the second equals

$$\sum_{i=n+1}^{m} \sum_{j=1}^{n} \left\langle \boldsymbol{E}^{ij}, \boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}^{ij}] f(\boldsymbol{\Sigma}) \right\rangle = \sum_{i=1}^{m-n} \sum_{j=1}^{n} \left\langle \boldsymbol{E}^{ij}, \boldsymbol{E}^{ij} \boldsymbol{\Sigma}^{-1} f(\boldsymbol{\Sigma}) \right\rangle = (m-n) \sum_{i=1}^{n} \frac{f_i(\sigma_i)}{\sigma_i}.$$

Since $n = \min(m,n)$, we conclude that

$$\operatorname{div}(f(\boldsymbol{X})) = S^{\tilde{\boldsymbol{U}}} + S^{\boldsymbol{\Sigma}} + S^{\boldsymbol{V}} = \sum_{i=1}^{\min(m,n)} \left(f_i'(\sigma_i) + |m-n| \frac{f_i'(\sigma_i)}{\sigma_i} \right) + 2 \sum_{i\neq j,\, i,j=1}^{\min(m,n)} \frac{\sigma_i f_i(\sigma_i)}{\sigma_i^2-\sigma_j^2}.$$

■

4.2 The complex case

Since the singular values are not analytic functions of the entries, we can only consider the derivatives of f as a function of the real and imaginary parts of the entries of $\boldsymbol{X}$. Therefore, we identify $\mathbb{C}^{m\times n}$ with $\mathbb{R}^{m\times n} \times \mathbb{R}^{m\times n}$ and f with a complex-valued function of $2mn$ variables. This has the particular consequence that the differential $\mathrm{d}f$ at $\boldsymbol{X} \in \mathbb{C}^{m\times n}$—whenever it exists—obeys

$$\mathrm{d}f_{\boldsymbol{X}}[\boldsymbol{\Delta}] = \mathrm{d}f_{\boldsymbol{X}}[\mathrm{Re}(\boldsymbol{\Delta})] + \mathrm{d}f_{\boldsymbol{X}}[\imath \mathrm{Im}(\boldsymbol{\Delta})],$$

for all $\boldsymbol{\Delta} \in \mathbb{C}^{m\times n}$; here, $\imath$ is the imaginary unit. The differential seen as a function with domain $\mathbb{R}^{m\times n} \times \mathbb{R}^{m\times n}$ is of course linear. Lemma 4.2 may be adapted to the complex setting with minor modifications.

Lemma 4.4. *Let $\boldsymbol{X}$ be simple and full-rank. Then the differentials of $\boldsymbol{U}$ and $\boldsymbol{V}$ are given by*

$$\mathrm{d}\boldsymbol{U}[\boldsymbol{\Delta}] = \tilde{\boldsymbol{U}}\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}], \qquad \mathrm{d}\boldsymbol{V}[\boldsymbol{\Delta}] = \boldsymbol{V}\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}]^*,$$

where $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}$ and $\boldsymbol{\Omega}_{\boldsymbol{V}}$ are given in closed-form (4.15) *($\mathrm{d}\boldsymbol{U}$ is independent of the choice of $\tilde{\boldsymbol{U}}$). Further, $\mathrm{d}\boldsymbol{\Sigma}$ is given by* (4.13)*. For completeness, $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}, \boldsymbol{\Omega}_{\boldsymbol{V}} : \mathbb{C}^{m\times n} \mapsto \mathbb{C}^{m\times n}$ while $\mathrm{d}\boldsymbol{\Sigma} : \mathbb{C}^{m\times n} \mapsto \mathbb{R}^{m\times n}$.*

Proof Following the same steps as in the proof of Lemma 4.2, let $\boldsymbol{X} \in \mathbb{C}^{m\times n}$ ($m \geq n$). We see that (4.4) becomes

$$\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V} = \boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{\Delta}]\boldsymbol{\Sigma} + \boldsymbol{I}_{m\times n}\mathrm{d}\boldsymbol{\Sigma}[\boldsymbol{\Delta}] + \boldsymbol{I}_{m\times n}\boldsymbol{\Sigma}\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}], \tag{4.12}$$

where $\boldsymbol{\Omega}_{\boldsymbol{V}}$ and the upper $n \times n$ block of $\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}$ are now skew-Hermitian. Since these matrices have purely imaginary entries on the diagonal,

$$\mathrm{d}\sigma_i[\boldsymbol{\Delta}] = \mathrm{Re}\{(\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ii}\} = \frac{(\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ii} + \overline{(\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ii}}}{2}, \qquad i = 1, \ldots, n, \tag{4.13}$$

and

$$\Omega_{\tilde{\boldsymbol{U}},ii}[\boldsymbol{\Delta}]\sigma_i + \Omega_{\boldsymbol{V},ii}[\boldsymbol{\Delta}]\sigma_i = \mathrm{Im}\{(\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ii}\}\imath = \frac{(\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ii} - \overline{(\tilde{\boldsymbol{U}}^*\boldsymbol{\Delta}\boldsymbol{V})_{ii}}}{2}, \qquad i = 1, \ldots, n. \tag{4.14}$$

Furthermore, it is not hard to see that the system (4.8) becomes

$$\begin{bmatrix}\Omega_{\tilde{\boldsymbol{U}},ij}[\boldsymbol{\Delta}] \\ \Omega_{\boldsymbol{V},ij}[\boldsymbol{\Delta}]\end{bmatrix} = -\frac{1}{\sigma_i^2 - \sigma_j^2}\begin{bmatrix}\sigma_j & \sigma_i \\ -\sigma_i & -\sigma_j\end{bmatrix}\begin{bmatrix}(\boldsymbol{U}^*\boldsymbol{\Delta}\boldsymbol{V})_{ij} \\ \overline{(\boldsymbol{U}^*\boldsymbol{\Delta}\boldsymbol{V})_{ji}}\end{bmatrix}, \qquad i \neq j, \quad i, j = 1, \ldots, n, \tag{4.15}$$

and the expression equivalent to (4.9) is

$$(\tilde{\boldsymbol{U}}^T\boldsymbol{\Delta}\boldsymbol{V})_{ij} = \Omega_{\tilde{\boldsymbol{U}},ij}[\boldsymbol{\Delta}]\sigma_j, \qquad i = n+1, \ldots, m, \quad j = 1, \ldots, n. \tag{4.16}$$

As in the real case, this shows $\mathrm{d}\boldsymbol{\Sigma}[\boldsymbol{\Delta}]$, $\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{\Delta}]$ and $\boldsymbol{\Omega}_{\boldsymbol{U}}[\boldsymbol{\Delta}]$ act linearly on $\boldsymbol{\Delta}$, and that they are uniquely characterized[3] by the SVD of $\boldsymbol{X}$. Furthermore, the same argument allows us to show that the differential of $\boldsymbol{U}$ is independent of the choice of $\tilde{\boldsymbol{U}}$, concluding the proof. ∎

As in the real case, we restrict attention to a simple and full-rank $\boldsymbol{X} \in \mathbb{C}^{m\times n}$. We need to determine two terms for the divergence as discussed in Section 3.2. Now if $\{\boldsymbol{\Delta}^{ij}\}_{(i,j)\in\mathcal{I}}$ is an orthonormal basis for $\mathbb{R}^{m\times n}$, then $\{\imath\boldsymbol{\Delta}^{ij}\}_{(i,j)\in\mathcal{I}}$ is an orthonormal basis for $\imath\mathbb{R}^{m\times n}$. With the divergence as in (3.6),

$$\begin{aligned}\mathrm{div}\,(f(\boldsymbol{X})) &= \sum_{(i,j)\in\mathcal{I}} \mathrm{Re}(\langle \boldsymbol{\Delta}^{ij}, \mathrm{d}f[\boldsymbol{\Delta}^{ij}]\rangle) + \sum_{(i,j)\in\mathcal{I}} \mathrm{Im}(\langle \boldsymbol{\Delta}^{ij}, \mathrm{d}f[\imath\boldsymbol{\Delta}^{ij}]\rangle) \\ &= \mathrm{Re}\left(\sum_{(i,j)\in\mathcal{I}} \langle \boldsymbol{\Delta}^{ij}, \mathrm{d}f[\boldsymbol{\Delta}^{ij}] - \imath\mathrm{d}f[\imath\boldsymbol{\Delta}^{ij}]\rangle\right).\end{aligned}$$

The analogue of Theorem 4.3 is this:

[3]The attentive reader will notice that (4.14) leaves us with one degree of freedom, and thus does not determine *exactly* the differentials. This is not a problem, and it is a consequence of the freedom of choice of a phase factor for the singular vectors.

Theorem 4.5. *Let f be a matrix-valued spectral function and suppose $\boldsymbol{X} \in \mathbb{C}^{m\times n}$ is simple and has full-rank. Then*

$$\operatorname{div}(f(\boldsymbol{X})) = \sum_{i=1}^{\min(m,n)} \left(f_i'(\sigma_i) + (2|m-n|+1)\frac{f_i(\sigma_i)}{\sigma_i} \right) + 4 \sum_{i\neq j,\, i,j=1}^{\min(m,n)} \frac{\sigma_i f_i(\sigma_i)}{\sigma_i^2 - \sigma_j^2}. \tag{4.17}$$

Proof Since $\boldsymbol{X}$ is simple and has full-rank, Lemma 4.4 holds, and

$$\begin{aligned} \mathrm{d}f[\boldsymbol{E}^{ij}] - \imath \mathrm{d}f[\imath\boldsymbol{E}^{ij}] = \tilde{\boldsymbol{U}} \left[(\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{E}^{ij}] - \imath\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\imath\boldsymbol{E}^{ij}]) f(\boldsymbol{\Sigma}) + \boldsymbol{I}_{m\times n} f(\boldsymbol{\Sigma})(\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{E}^{ij}] - \imath\boldsymbol{\Omega}_{\boldsymbol{V}}[\imath\boldsymbol{E}^{ij}]) \right] \boldsymbol{V}^* \\ + \tilde{\boldsymbol{U}} \boldsymbol{I}_{m\times n} \mathrm{d}f_{\boldsymbol{\Sigma}}[\mathrm{d}\boldsymbol{\Sigma}[\boldsymbol{E}^{ij}] - \imath \mathrm{d}\boldsymbol{\Sigma}[\imath\boldsymbol{E}^{ij}]] \boldsymbol{V}^*. \end{aligned} \tag{4.18}$$

Put $\boldsymbol{L}^{ij} = \tilde{\boldsymbol{U}}^* \boldsymbol{E}^{ij} \boldsymbol{V}$ for $(i,j) \in \mathcal{I}$ for short. From (4.13), we get

$$\mathrm{d}\sigma_i[\boldsymbol{E}^{ij}] - \imath \mathrm{d}\sigma_i[\imath\boldsymbol{E}^{ij}] = L_{kk}^{ij}$$

for $(i,j) \in \mathcal{I}$ and $k = 1, \ldots, n$. By the same arguments, (4.15) gives

$$\begin{bmatrix} \Omega_{\tilde{\boldsymbol{U}},kl}[\boldsymbol{E}^{ij}] - \imath\Omega_{\tilde{\boldsymbol{U}},kl}[\imath\boldsymbol{E}^{ij}] \\ \Omega_{\boldsymbol{V},kl}[\boldsymbol{E}^{ij}] - \imath\Omega_{\boldsymbol{V},kl}[\boldsymbol{E}^{ij}] \end{bmatrix} = -\frac{1}{\sigma_k^2 - \sigma_l^2} \begin{bmatrix} \sigma_l & \sigma_k \\ -\sigma_k & -\sigma_l \end{bmatrix} \begin{bmatrix} 2L_{kl}^{ij} \\ 0 \end{bmatrix},$$

for $(i,j) \in \mathcal{I}$, $k, l = 1, \ldots, n$, and $k \neq l$. Next, (4.16) gives

$$\Omega_{\tilde{\boldsymbol{U}},kl}[\boldsymbol{E}^{ij}] = 2\frac{L_{kl}^{ij}}{\sigma_l},$$

for $(i,j) \in \mathcal{I}$, $k = n+1, \ldots, m$, and $l = 1, \ldots, n$. Finally, (4.14) implies

$$\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}},kk}[\boldsymbol{E}^{ij}] + \Omega_{\boldsymbol{V},kk}[\boldsymbol{E}^{ij}] - \imath(\Omega_{\tilde{\boldsymbol{U}},kk}[\imath\boldsymbol{E}^{ij}] + \Omega_{\boldsymbol{V},kk}[\imath\boldsymbol{E}^{ij}]) = \frac{L_{kk}^{ij}}{\sigma_k},$$

for $k = 1, \ldots, n$. Now, set

$$\begin{aligned} S^{\tilde{\boldsymbol{U}}} &= \sum_{(i,j)\in\mathcal{I}} \left\langle \boldsymbol{L}^{ij}, (\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\boldsymbol{E}^{ij}] - \imath\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}}}[\imath\boldsymbol{E}^{ij}]) f(\boldsymbol{\Sigma}) \right\rangle, \\ S^{\boldsymbol{\Sigma}} &= \sum_{(i,j)\in\mathcal{I}} \left\langle \boldsymbol{L}^{ij}, \boldsymbol{I}_{m\times n} f(\boldsymbol{\Sigma})(\boldsymbol{\Omega}_{\boldsymbol{V}}[\boldsymbol{E}^{ij}] - \imath\boldsymbol{\Omega}_{\boldsymbol{V}}[\imath\boldsymbol{E}^{ij}]) \right\rangle, \\ S^{\boldsymbol{V}} &= \sum_{(i,j)\in\mathcal{I}} \left\langle \boldsymbol{L}^{ij}, \boldsymbol{I}_{m\times n} \mathrm{d}f_{\boldsymbol{\Sigma}}[\mathrm{d}\boldsymbol{\Sigma}[\boldsymbol{E}^{ij}] - \imath\mathrm{d}\boldsymbol{\Sigma}[\imath\boldsymbol{E}^{ij}]] \right\rangle. \end{aligned}$$

Then it is clear that

$$\begin{aligned} S^{\boldsymbol{V}} &= \sum_{(i,j)\in\mathcal{I}} \sum_{k\neq l,\, k,l=1}^{n} \frac{2f_k(\sigma_k)\sigma_k |L_{kl}^{ij}|^2}{\sigma_k^2 - \sigma_l^2} + \sum_{k=1}^{n} f_k(\sigma_k) \sum_{(i,j)\in\mathcal{I}} \overline{L_{kk}^{ij}} (\boldsymbol{\Omega}_{\boldsymbol{V},kk}[\boldsymbol{E}^{ij}] - \imath\boldsymbol{\Omega}_{\boldsymbol{V},kk}[\imath\boldsymbol{E}^{ij}]) \\ &= 2 \sum_{k\neq l,\, k,l=1}^{n} \frac{f_k(\sigma_k)\sigma_k}{\sigma_k^2 - \sigma_l^2} + \sum_{k=1}^{n} f_k(\sigma_k) \sum_{(i,j)\in\mathcal{I}} \overline{L_{kk}^{ij}} (\boldsymbol{\Omega}_{\boldsymbol{V},kk}[\boldsymbol{E}^{ij}] - \imath\boldsymbol{\Omega}_{\boldsymbol{V},kk}[\imath\boldsymbol{E}^{ij}]), \end{aligned}$$

and

$$S^{\boldsymbol{\Sigma}} = \sum_{(i,j)\in\mathcal{I}} \sum_{k=1}^{n} f_k'(\sigma_k)|L_{kk}^{ij}|^2 = \sum_{k=1}^{n} f_k'(\sigma_k).$$

Finally,

$$\begin{aligned} S^{\tilde{\boldsymbol{U}}} = &- \sum_{(i,j)\in\mathcal{I}} \sum_{k\neq l,\, k,l=1}^{n} \frac{2f_l(\sigma_l)\sigma_l|L_{kl}^{ij}|^2}{\sigma_k^2-\sigma_l^2} + \sum_{k=1}^{n} f_k(\sigma_k) \sum_{(i,j)\in\mathcal{I}} \overline{L_{kk}^{ij}}(\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}},kk}[\boldsymbol{E}^{ij}] - \imath\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}},kk}[\imath\boldsymbol{E}^{ij}]) \\ &+ \sum_{(i,j)\in\mathcal{I}} \sum_{k=n+1}^{m} \sum_{l=1}^{n} 2\frac{f_l(\sigma_l)|L_{kl}^{ij}|^2}{\sigma_l} \\ = 2 &\sum_{k\neq l,\, k,l=1}^{n} \frac{f_k(\sigma_k)\sigma_k}{\sigma_k^2-\sigma_l^2} + 2(m-n)\sum_{k=1}^{n}\frac{f_k(\sigma_k)}{\sigma_k} + \sum_{k=1}^{n} f_k(\sigma_k) \sum_{(i,j)\in\mathcal{I}} \overline{L_{kk}^{ij}}(\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}},kk}[\boldsymbol{E}^{ij}] - \imath\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}},kk}[\imath\boldsymbol{E}^{ij}]). \end{aligned}$$

Since

$$\sum_{(i,j)\in\mathcal{I}} \overline{L_{kk}^{ij}} \left(\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}},kk}[\boldsymbol{E}^{ij}] - \imath\boldsymbol{\Omega}_{\tilde{\boldsymbol{U}},kk}[\imath\boldsymbol{E}^{ij}] + \boldsymbol{\Omega}_{\boldsymbol{V},kk}[\boldsymbol{E}^{ij}] - \imath\boldsymbol{\Omega}_{\boldsymbol{V},kk}[\imath\boldsymbol{E}^{ij}]\right) = \sum_{(i,j)\in\mathcal{I}} \frac{|L_{kk}^{ij}|^2}{\sigma_k} = \frac{1}{\sigma_k},$$

we conclude

$$\begin{aligned} \operatorname{div}(f(\boldsymbol{X})) &= \operatorname{Re}\left(S^{\tilde{\boldsymbol{U}}} + S^{\boldsymbol{\Sigma}} + S^{\boldsymbol{V}}\right) \\ &= \sum_{i=1}^{\min(m,n)} \left(f_i'(\sigma_i) + (2|m-n|+1)\frac{f_i(\sigma_i)}{\sigma_i}\right) + 4\sum_{i\neq j,\, i,j=1}^{\min(m,n)} \frac{\sigma_i f_i(\sigma_i)}{\sigma_i^2-\sigma_j^2}, \end{aligned}$$

where we used $n = \min(m,n)$. ∎

4.3 An extension to all matrix arguments

Theorems 4.3 and 4.5 provide the divergence of a matrix-valued spectral function in closed-form for simple, full-rank arguments. Here, we indicate that these formulas have a continuous extension to all matrices. Before continuing, we introduce some notation. For a given $\boldsymbol{X}$, we denote as $\kappa \leq n$ the number of distinct singular values while s_i and d_i, $i = 1, \ldots, \kappa$, denote the i-th distinct singular value and its multiplicity.

As previously discussed, when $\boldsymbol{\Sigma}$ is simple there is no ambiguity as to which f_i is being applied to a singular value. However, when $\boldsymbol{\Sigma}$ is not simple, the ambiguities make f nondifferentiable unless $f_i \equiv f$ for $i = 1, \ldots, n$.

Theorem 4.6. *Let f be a matrix-valued spectral function, with $f_i \equiv f$ for $i = 1, \ldots, n$, and $f(0) = 0$. Then the divergence extends by continuity to any $\boldsymbol{X} \in \mathbb{R}^{m\times n}$, and is given by*

$$\begin{aligned} \operatorname{div}(f(\boldsymbol{X})) &= \sum_{i:\, s_i>0} \left[\left(d_i + \binom{d_i}{2}\right) f'(s_i) + \left(|m-n|d_i + \binom{d_i}{2}\right)\frac{f(s_i)}{s_i}\right] \\ &+ \sum_{i:\, s_i=0} \left[(|m-n|+1)d_i + 2\binom{d_i}{2}\right] f'(0) + 2\sum_{i\neq j,\, i,j=1}^{\kappa} d_i d_j \frac{s_i f(s_i)}{s_i^2 - s_j^2}, \end{aligned} \tag{4.19}$$

and to any $\boldsymbol{X} \in \mathbb{C}^{m\times n}$,

$$\begin{aligned}\operatorname{div}(f(\boldsymbol{X})) &= \sum_{i:\, s_i>0}\left[\left(d_i + 2\binom{d_i}{2}\right) f'(s_i) + \left((2|m-n|+1)d_i + 2\binom{d_i}{2}\right)\frac{f(s_i)}{s_i}\right] \\ &+ \sum_{i:\, s_i=0}\left[2(|m-n|+1)d_i + 4\binom{d_i}{2}\right] f'(0) + 4\sum_{i\neq j,\, i,j=1}^{\kappa} d_i d_j \frac{s_i f(s_i)}{s_i^2 - s_j^2}. \qquad (4.20)\end{aligned}$$

Proof We only prove (4.19) as the same arguments apply to (4.20). The key idea involves non-tangential limits defined below. Suppose we have a sequence $\{\boldsymbol{X}_k\}$ of full-rank matrices with $\boldsymbol{X}_k \to \boldsymbol{X}$ as $k \to \infty$; we establish that the limit $\operatorname{div}(f(\boldsymbol{X}_k))$ is independent of the sequence. Since (4.10) depends only on the singular values of $\boldsymbol{X}_k$, we can focus on the sequence of singular values. Let $\{\sigma_i\}_{i=1}^n$ be the singular values of $\boldsymbol{X}$, and let $\{\sigma_i^k\}_{i=1}^n$ be those of $\boldsymbol{X}_k$. We have

$$\operatorname{div}(f(\boldsymbol{X}_k)) = \sum_{i=1}^{\min(m,n)}\left(f'(\sigma_i^k) + |m-n|\frac{f(\sigma_i^k)}{\sigma_i^k}\right) + 2\sum_{i\neq j,\, i,j=1}^{\min(m,n)} \frac{\sigma_i^k f(\sigma_i^k)}{(\sigma_i^k)^2 - (\sigma_j^k)^2}.$$

The first terms are easy to analyze. In fact,

$$\lim_{k\to\infty} f'(\sigma_i^k) = f'(\sigma_i) \quad \text{and} \quad \lim_{k\to\infty}\frac{f(\sigma_i^k)}{\sigma_i^k} = \begin{cases}\frac{f(\sigma_i)}{\sigma_i}, & \text{if } \sigma_i > 0,\\ f'(0), & \text{if } \sigma_i = 0.\end{cases}$$

The remaining term can be decomposed as

$$\sum_{i\neq j,\, i,j=1}^{\min(m,n)} \frac{\sigma_i^k f(\sigma_i^k)}{(\sigma_i^k)^2 - (\sigma_j^k)^2} = \left[\sum_{i\neq j,\, \sigma_i=\sigma_j}^{\min(m,n)} + \sum_{i\neq j,\, \sigma_i\neq\sigma_j}^{\min(m,n)}\right] \frac{\sigma_i^k f(\sigma_i^k)}{(\sigma_i^k)^2 - (\sigma_j^k)^2} := I_0 + I_1. \qquad (4.21)$$

We now formalize the notion of non-tangential limit. Consider a pair (i,j) with $i \neq j$ and $\sigma_i = \sigma_j := \sigma$, and define

$$\delta_i^k = \sigma_i^k - \sigma, \quad \delta_j^k = \sigma_j^k - \sigma \quad \text{and} \quad \rho_{ij}^k = \delta_j^k/\delta_i^k.$$

We say that $\boldsymbol{X}_k$ approaches $\boldsymbol{X}$ non-tangentially if there exists a constant $c_{ij} > 0$ such that $|\rho_{ij}^k - 1| > c_{ij}$. Roughly speaking, the matrices $\boldsymbol{X}_k$ approach $\boldsymbol{X}$ fast enough, so that they remain simple, and have full-rank. We have

$$\begin{aligned}\frac{\sigma_i^k f(\sigma_i^k) - \sigma_j^k f(\sigma_j^k)}{(\sigma_i^k)^2 - (\sigma_j^k)^2} &= \frac{\sigma}{2\sigma + \delta_i^k + \delta_j^k}\left(\frac{f(\sigma+\delta_i^k) - f(\sigma+\delta_j^k)}{\delta_i^k - \delta_j^k}\right) \\ &+ \frac{1}{2\sigma + \delta_i^k + \delta_j^k}\left(\frac{1}{1-\rho_{ij}^k} f(\sigma+\delta_i^k) - \frac{\rho_{ij}^k}{1-\rho_{ij}^k} f(\sigma+\delta_j^k)\right).\end{aligned}$$

We distinguish two cases. If $\sigma > 0$, we see that

$$\lim_{k\to\infty}\frac{\sigma_i^k f(\sigma_i^k) - \sigma_j^k f(\sigma_j^k)}{(\sigma_i^k)^2 - (\sigma_j^k)^2} = \frac{1}{2}f'(\sigma) + \frac{1}{2}\frac{f(\sigma)}{\sigma}$$

whereas when $\sigma = 0$ we obtain

$$\lim_{k\to\infty} \frac{\sigma_i^k f(\sigma_i^k) - \sigma_j^k f(\sigma_j^k)}{(\sigma_i^k)^2 - (\sigma_j^k)^2} = f'(0)$$

and, therefore,

$$I_0 = \sum_{i:\, s_i=0} \binom{d_i}{2} f'(0) + \frac{1}{2} \sum_{i:\, s_i>0} \binom{d_i}{2} \left(f'(s_i) + \frac{f(s_i)}{s_i} \right).$$

The second sum I_1 in (4.21) is easier to analyze, as it can be seen that

$$\lim_{k\to\infty} \sum_{i\neq j,\, \sigma_i\neq\sigma_j}^{\min(m,n)} \frac{\sigma_i^k f(\sigma_i^k)}{(\sigma_i^k)^2 - (\sigma_j^k)^2} = \sum_{i\neq j,\, i,j=1}^{\kappa} d_i d_j \frac{s_i f(s_i)}{s_i^2 - s_j^2},$$

concluding the proof. ■

5 Discussion

Beyond MRI, SVT denoising also has potential application for medical imaging modalities outside of MRI. For example, in perfusion X-ray computed tomography (CT) studies, where some anatomical region (e.g., the kidneys) is repeatedly scanned to visualize hemodynamics, X-ray tube energy is routinely lowered to limit the ionizing radiation dose delivered to the patient. However, this can result in a substantial increase of image noise. Recent works (e.g., [4]) have shown that retrospective denoising of low-dose CT data can often yield diagnostic information comparable to full-dose images. It may be possible to adapt SVT, and the proposed risk estimators, for the CT dose reduction problem as well.

Following the successful application of compressive sensing theory [7, 9] for accelerated MRI (e.g., [20, 35]), the use of low-rank constraints for undersampled MRI reconstruction is also of increasing interest. For the dynamic and parametric acquisition scenarios described above, the generalization from denoising to undersampled reconstruction is straightforward in light of recent developments on the topic of matrix completion [6, 29]. To date, low-rank reconstruction methods have been deployed for accelerated cardiac imaging [13, 19, 36] and dynamic contrast-enhanced MRI (DCE-MRI) of the breast [13]. More recently, it has been demonstrated that reconstructing accelerated parallel MRI data under low-rank constraints [21, 37] can obviate the need for expensive and potentially error-prone patient-specific calibration procedures as well as facilitate previously impractical acquisition strategies. As with the denoising application, parameter selection is an outstanding challenge in undersampled reconstruction problems. Unlike the denoising problem, only incomplete noisy data is available and Stein's method is not directly applicable. Recently, for undersampled MRI reconstruction, Ramani et al. [28] demonstrated that Stein's method can be used to generate an unbiased risk estimator for the predicted MSE of ℓ_1-minimization reconstructions, and that this can be used for automated parameter optimization. An interesting future direction of research will lie in determining if the models developed in this work facilitate development of a similar framework for the low-rank matrix recovery problem.

Acknowledgements

E. C. is partially supported by AFOSR under grant FA9550-09-1-0643, by ONR under grant N00014-09-1-0258 and by a gift from the Broadcom Foundation. C. S-L. is partially supported by a Fulbright-CONICYT Scholarship and by the Simons Foundation. J. D. T. and the Mayo Clinic Center for Advanced Imaging Research are partially supported by National Institute of Health grant RR018898. E. C. would like to thank Rahul Mazumder for fruitful discussions about this project.

References

[1] H. H. Bauschke and P. L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. CMS Books in Mathematics. Springer, 2011.

[2] T. Blu and F. Luisier. The SURE-LET approach to image denoising. *IEEE Transactions on Image Processing*, 16(11):2778–2786, 2007.

[3] J. Bogaert, S. Dymarkowski, and A. M. Taylor. *Clinical cardiac MRI*. Medical Radiology: Diagnostic Imaging. Springer, 2005.

[4] A. Borsdorf, R. Raupach, T. Flohr, and J. Hornegger. Wavelet based noise reduction in CT-images using correlation analysis. *IEEE Transactions on Medical Imaging*, 27(12):1695–1703, 2008.

[5] J. A. Cadzow. Signal enhancement - a composite property mapping algorithm. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 36(1):49–62, 1988.

[6] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

[7] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.

[8] C.-A. Deledalle, S. Vaiter, G. Peyré, J. Fadili, and C. Dossal. Risk estimation for matrix recovery with spectral regularization. arXiv:1205.1482v1, May 2012.

[9] D.L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.

[10] A. Edelman. Handouts for 18.325: Matrix jacobians without wedge products, 2005.

[11] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Annals of Statistics*, 32:407–499, 2004.

[12] L. C. Evans and R. F. Gariepy. *Measure theory and fine properties of functions*. Boca Raton, FL: CRC Press, 1992.

[13] J. Haldar and Z.-P. Liang. spatiotemporal imaging with partially separable functions: a matrix recovery approach. In *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 716–719, 2010.

[14] I. Johnstone. *Gaussian estimation: Sequence and wavelet models*. Manuscript, December 2011.

[15] F. Lam, S. D. Babacan, J. P. Haldar, N. Schuff, and Z.-P. Liang. Denoising diffusion-weighted MR magnitude image sequences using low rank and edge constraints. In *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1401–1404, 2012.

[16] A. Lewis and H. Sendov. Twice differentiable spectral functions. *SIAM Journal on Matrix Analysis and Applications*, 23(2):368–386, 2001.

[17] Z.-P. Liang. Spatiotemporal imaging with partially separable functions. In *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 988–991, 2007.

[18] Z.-P. Liang and P. C. Lauterbur. *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. Biomedical Engineering. IEEE, 2000.

[19] S. G. Lingala, Y. Hu, E. DiBella, and M. Jacob. Accelerated dynamic MRI exploiting sparsity and low-rank structure: kt-SLR. *IEEE Transactions on Medical Imaging*, 30(5):1042–1054, 2011.

[20] M. Lustig, D. L. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.

[21] M. Lustig, M. Elad, and J. M. Pauly. Calibrationless parallel imaging reconstruction by structured low-rank matrix completion. In *Proceedings of the 18th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*, page 2870, 2010.

[22] Bydder M. and Du J. Noise reduction in multiple-echo data sets using singular value decomposition. *Magnetic Resonance Imaging*, 24(7):849–856, 2006.

[23] J. R. Magnus. On differentiating eigenvalues and eigenvectors. *Econometric Theory*, 1(2):pp. 179–191, 1985.

[24] J. V. Manjon, P. Coupe, L. Concha, A. Buades, L. Collins, and M. Robles. DWI denoising using overcomplete local PCA decomposition. In *Proceedings of the 20th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*, page 3580, 2012.

[25] H. M. Nguyen, X. Peng, M. N. Do, and Z.-P. Liang. Spatiotemporal denoising of MR spectroscopic imaging data by low-rank approximations. In *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 857–860, 2011.

[26] J. N. Oshinski, J. G. Delfino, P. Sharma, A. M. Gharib, and Pettigrew R. I. Cardiovascular magnetic resonance at 3.0 T: current state of the art. *Journal of Cardiac Magnetic Resonance*, 12(1):55, 2010.

[27] T. Papadopoulo and M. I. A. Lourakis. Estimating the jacobian of the singular value decomposition: Theory and applications. In *Proceedings of the European Conference on Computer Vision, ECCV00*, pages 554–570. Springer, 2000.

[28] S. Ramani, Z. Liu, J. Rosen, J.-F. Nielsen, and J. A. Fessler. Regularization parameter selection for nonlinear iterative image restoration and MRI reconstruction using GCV and SURE-based methods. *IEEE Transactions on Image Processing*, 21(8):3659–3672, 2012.

[29] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum rank solutions to linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.

[30] R. T. Rockafellar. *Convex analysis*. Princeton landmarks in mathematics. Princeton, NJ: Princeton University Press, 1997.

[31] B. Shariff and Y. Bresler. Physiologically improved NCAT Phantom (PINCAT) enables in-silico study of the effects of beat-to-beat variability on cardiac MR. In *Proceedings of the 15th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*, page 3418, 2007.

[32] C. M. Stein. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, 9(6):pp. 1135–1151, 1981.

[33] D. Sun and J. Sun. Nonsmooth matrix valued functions defined by singular values. http://www.math.nus.edu.sg/ matsundf/SS3.pdf, December 2002.

[34] C. G. Thomas, R. A. Harshman, and R. S. Menon. Noise reduction in BOLD-based fMRI using component analysis. *NeuroImage*, 17(3):1521–1537, 2002.

[35] J. D. Trzasko and A. Manduca. Highly undersampled magnetic resonance image reconstruction via homotopic ℓ_0 -minimization. *IEEE Transactions on Medical Imaging*, 28(1):106–121, 2009.

[36] J. D. Trzasko and A. Manduca. Local versus global low-rank promotion in dynamic MRI series reconstruction. In *Proceedings of the 19th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*, page 4371, 2011.

[37] J. D. Trzasko and A. Manduca. CLEAR: Calibration-free parallel imaging using locally low-rank encouraging reconstruction. In *Proceedings of the 20th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*, page 517, 2012.