

Oden Institute REPORT 202024

December 2020

Data-Informed Regularization For Inverse and Imaging Problems

by

Jonathan Wittmer and Tan Bui-Thanh



Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin
Austin, Texas 78712

Reference: Jonathan Wittmer and Tan Bui-Thanh, "Data-Informed Regularization For Inverse and Imaging Problems," Oden Institute REPORT 20-24, Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, December 2020.

DATA-INFORMED REGULARIZATION FOR INVERSE AND IMAGING PROBLEMS*

JONATHAN WITTMER[†] AND TAN BUI-THANH[‡]

1. Introduction. Regularization is a technique often employed to facilitate the well-posedness of inverse (and imaging) problems. An inverse solution is thus a trade-off between the data misfit and the regularization. Due to noise and limited availability, the data typically informs limited directions in the parameter space where the inverse solution resides. A desired regularization, we argue, should minimally interfere with these data-informed directions. However, most regularization techniques regularize all parameter directions, including the data-informed ones, thus polluting the resulting inverse solution. This chapter presents a new regularization method for inverse and imaging problems, called data-informed regularization, that implicitly avoids regularizing the data-informed directions. The DI method combines advantages of the classical truncated SVD and Tikhonov regularization. In particular, it does not pollute the data-informed modes with regularization, but regularizes only the less data-informed ones.

Compared to existing approaches, our method has many distinct and advantageous features: 1) it automatically determines the directions equally informed by the data and any Tikhonov regularization while leaving the most informative directions untouched. In fact, we will show that, similar to the balanced truncation idea in control theory (see, e.g., [4, 1] and the references therein), this is done implicitly by seeking directions in parameter space that balance the information from regularization and data, and removing the regularization on them; 2) We will show that our approach has an intuitive statistical interpretation, namely, it transforms both the data distribution (i.e. the likelihood) and prior distribution (induced by Tikhonov regularization) to the same Gaussian distribution whose covariance matrix is diagonal and the diagonal elements are exactly the singular values of a composition of the prior covariance matrix, the forward map, and the noise covariance matrix. 3) Though constructively derived and its insights obtained from the truncated Singular Value Decomposition (SVD), the inverse solution resulting from our approach does not necessarily require an SVD decomposition, which may not be feasible for large-scale applications. Indeed, we will present a nested matrix-free approach to obtain an approximate inverse solution. 4) By construction, features in our inverse solution dictated by the data-informed directions are insensitive to the regularization parameter. For many inverse and imaging problems, these features dominate the solution and thus the inverse solution resulting from our regularization technique is robust with respect to regularization parameter values. These findings will be demonstrated and supported by various numerical results from deblurring, denoising, and X-ray tomography problems.

*Submitted to the editors DATE.

Funding: This work was partially funded by the National Science Foundation awards NSF-1808576 and NSF-CAREER-1845799; by the Defense Threat Reduction Agency award DTRA-M1802962; by the Department of Energy award DE-SC0018147; by KAUST; by 2018 ConTex award; and by 2018 UT-Portugal CoLab award. The authors are grateful to the supports.

[†]Department of Aerospace Engineering and Engineering Mechanics, UT Austin, Austin, Texas

[‡]Department of Aerospace Engineering and Engineering Mechanics, The Oden Institute for Computational Engineering and Sciences, UT Austin, Austin, Texas (tanbui@ices.utexas.edu, <https://users.oden.utexas.edu/~tanbui/>).

2. A data-informed regularization (DI) approach.

2.1. Data-informed regularization derivation. In this section we review the key ideas behind regularization by truncation using the Singular Value Decomposition (SVD). This provides the basic insights into the data-informed regularization technique. A statistical interpretation of the data-informed inverse framework will be discussed in Section 2.2. To begin, let us consider a linear inverse problem to determine $\mathbf{x} \in \mathbb{R}^p$ given

$$(2.1) \quad \mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e},$$

where $\mathbf{A} \in \mathbb{R}^{d \times p}$, $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \lambda^2 \mathbf{I})$, $\mathbf{I} \in \mathbb{R}^{d \times d}$ and $\mathbf{y} \in \mathbb{R}^d$. In the following, the identity matrix \mathbf{I} may have different size at different places and the actual size should be clear from the context. The simplest approach to attempt to solve this inverse problem is perhaps the least squares approach:

$$(2.2) \quad \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2,$$

where $\|\cdot\|$ denotes the standard Euclidean norm. The least-squares solution is given by

$$\mathbf{x}_{LS} = \left(\mathbf{A}^T \mathbf{A} \right)^{-1} \mathbf{A}^T \mathbf{y}.$$

Figure 1(a) plots the exact synthetic solution (black curve) against the least-squares solution (red curve) for a deconvolution problem with $d = p = 101$ and $\lambda = 0.05$. As can be seen, the least-squares solution blows up (or is unstable), which is not surprising since the inverse problem is ill-posed.

A standard Tikhonov regularization approach casts the above inverse problem into

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2 + \frac{\alpha}{2} \|\mathbf{x} - \mathbf{x}_0\|^2,$$

where \mathbf{x}_0 is given. A Tikhonov solution is presented in Figure 1(b) for $\alpha = 1$ and $\mathbf{x}_0 = \mathbf{0}$. Though this approach stabilizes the solution, it also smooths out the solution everywhere.

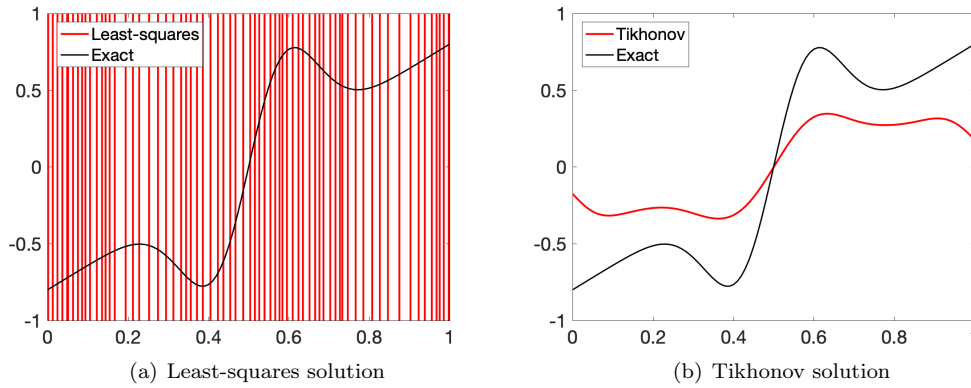


FIG. 1. Deconvolution using a) the least-squares approach, and b) a Tikhonov regularization with regularization parameter $\alpha = 1$ and $\mathbf{x}_0 = \mathbf{0}$.

Regularization by truncation does not require an explicit introduction of regularization term as in Tikhonov regularization. For example, the truncated SVD starts with the SVD decomposition of \mathbf{A} and then truncates all the singular vectors \mathbf{U}_j and \mathbf{V}_j corresponding to zero (or relatively small) singular values, i.e.,

$$\begin{aligned} \mathbf{A} &= \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T \\ &= \begin{pmatrix} \mathbf{U}_1 & \mathbf{U}_{n+1} & \mathbf{U}_d \\ \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots \end{pmatrix} \begin{pmatrix} \sigma_1 & & & \\ & \ddots & & \\ & & \sigma_n & \\ & & & 0 \\ & & & & \ddots \\ & & & & & 0 \end{pmatrix} \begin{pmatrix} \mathbf{V}_1^T \\ \vdots \\ \mathbf{V}_r^T \\ \mathbf{V}_{n+1}^T \\ \vdots \\ \mathbf{V}_p^T \end{pmatrix} \\ &= \mathbf{U}^n \mathbf{\Sigma}^n (\mathbf{V}^n)^T, \end{aligned}$$

where $\mathbf{U}^n := [\mathbf{U}_1, \dots, \mathbf{U}_n]$ (the first n columns of \mathbf{U} corresponding to n nonzero singular values or the range of the column space of \mathbf{A}), $\mathbf{\Sigma}^n := \text{diag}[\sigma_1, \dots, \sigma_n]$ ($\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$), $n \leq \min\{d, p\}$, and $\mathbf{V}^n := [\mathbf{V}_1, \dots, \mathbf{V}_n]$ (the first n columns of \mathbf{V} corresponding to n nonzero singular values the row space of \mathbf{A}).

The solution of the least-squares problem (2.2) using the truncated SVD together with pseudo-inverse reads

$$\mathbf{x}_{\text{SVD}}^n = (\mathbf{A}^T \mathbf{A})^\dagger \mathbf{A}^T \mathbf{y} = \mathbf{V}^n (\mathbf{\Sigma}^n)^{-1} (\mathbf{U}^n)^T \mathbf{y} = \sum_{i=1}^n \frac{\mathbf{U}_i^T \mathbf{y}}{\sigma_i} \mathbf{V}_i,$$

and

$$(2.3) \quad \mathbf{x}_{\text{SVD}}^r := \sum_{i=1}^r \frac{\mathbf{U}_i^T \mathbf{y}}{\sigma_i} \mathbf{V}_i,$$

which is a truncated SVD solution using r largest singular values. Figure 2 applies the truncated SVD approach to the deconvolution problem and compares the results with the Tikhonov regularization. As can be seen, truncated SVD solutions are stable and do not seem over-regularize the solution. However, as r increases, truncated SVD solutions tend to be more oscillatory (more unstable) because more “high-frequency” singular modes (corresponding to smaller singular values) are added to the truncated solutions. In other words, *the truncated SVD approach truncates the most unstable modes out of the solution. Why are these high-frequency modes undesirable?* The answer lies on the fact that the j th column of \mathbf{A} is the observational vector when the parameter \mathbf{x} is the j th canonical basis vector in \mathbb{R}^p . Thus the range space (column space) of \mathbf{A} is the observable subspace in \mathbb{R}^d . Within this observable subspace, the subspace spanned by the singular vector \mathbf{U}_j corresponding to larger (largest) singular value is more (most) observable whereas those corresponding to smaller (zero) singular values are less (not) observable. The truncated solution (2.3) clearly shows that the component of the data vector \mathbf{y} in the less (or non-) observable subspace, i.e. the product $\mathbf{U}_j^T \mathbf{y}$ (due to noise or error in the measurements), is amplified by $1/\sigma_j$, thus promoting an oscillatory or unstable solution.

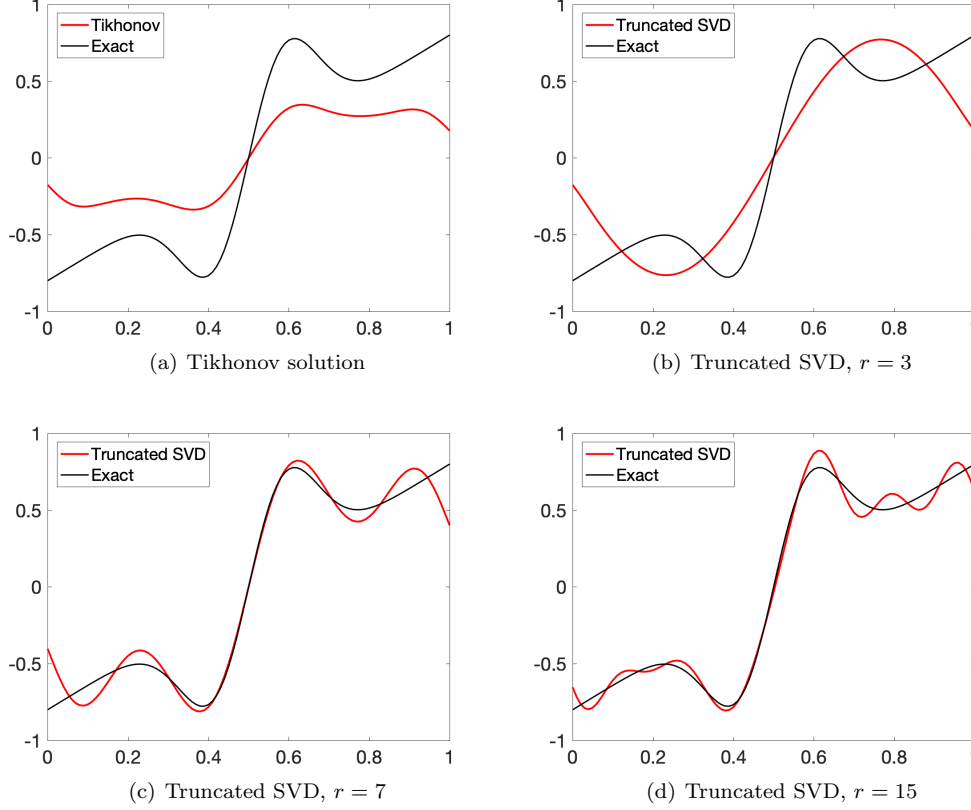


FIG. 2. Deconvolution using a) a Tikhonov regularization with regularization parameter $\alpha = 1$ and $\mathbf{x}_0 = \mathbf{0}$; b) truncated SVD with $r = 3$; c) truncated SVD with $r = 10$; and d) truncated SVD with $r = 15$.

In this chapter, we refer to row subspaces spanned by \mathbf{V}_j , corresponding to the **observable** subspaces spanned by \mathbf{U}_j , as **data-informed** parameter subspaces. Similarly, row subspaces spanned by \mathbf{V}_j , corresponding to the least observable subspaces spanned by \mathbf{U}_j , as **data-uninformed** parameter subspaces. The truncated SVD solution (2.3) clearly resides in the **data-informed** parameter subspaces for small r . The question is *where to truncate so that the solution is data-informed?*. The above discussion suggests that one should take r neither too large nor too small. Popular methods for picking r include the Morozov discrepancy principle, L-curve, and cross-validation, to name a few.

A closer look at the truncated SVD solution (2.3) shows that the truncated SVD approach zeros out the data-uninformed modes \mathbf{V}_j for $j \geq r + 1$. We next show that this is equivalent to infinitely regularizing data-uninformed directions. To see this, let us now consider a regularization scheme where the data-uninformed modes are penalized infinitely, i.e.,

$$(2.4) \quad \min \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2 + \frac{1}{2} \|\mathbf{L}(\mathbf{x} - \mathbf{x}_0)\|^2,$$

where

$$\begin{aligned} L^T L &:= \infty \left[\mathbf{I} - \mathbf{V}^r (\mathbf{V}^r)^T \right] = \infty (\mathbf{V}^r)^\perp \left((\mathbf{V}^r)^\perp \right)^T \\ &= \left[\mathbf{V}^r, (\mathbf{V}^r)^\perp \right] \left[\begin{array}{c|c} 0 & 0 \\ \hline 0 & \infty \mathbf{I} \end{array} \right] \left[\mathbf{V}^r, (\mathbf{V}^r)^\perp \right]^T, \end{aligned}$$

and $\left[\mathbf{I} - \mathbf{V}^r (\mathbf{V}^r)^T \right]$ is the orthogonal projection into the data-uninformed subspace spanned by $\{\mathbf{V}_j\}_{j=r+1}^d$. The solution of (2.4) is given by

$$\begin{aligned} \mathbf{x}_{Inf} &= \left\{ \mathbf{A}^T \mathbf{A} + \infty \left(\mathbf{I} - \mathbf{V}^r (\mathbf{V}^r)^T \right) \right\}^{-1} \mathbf{A}^T \mathbf{y} \\ &= \left\{ \left[\mathbf{V}^r, (\mathbf{V}^r)^\perp \right] \left(\left[\begin{array}{c|c} \Sigma_r^2 & 0 \\ \hline 0 & \mathbf{D} \end{array} \right] + \left[\begin{array}{c|c} 0 & 0 \\ \hline 0 & \infty \mathbf{I} \end{array} \right] \right) \left[\mathbf{V}^r, (\mathbf{V}^r)^\perp \right]^T \right\}^{-1} \mathbf{A}^T \mathbf{y} \\ &= \mathbf{V}^r (\Sigma^r)^{-2} (\mathbf{V}^r)^T \mathbf{A}^T \mathbf{y} = \mathbf{V}^r (\Sigma^r)^{-1} (\mathbf{U}^r)^T \mathbf{y} =: \mathbf{x}_{\text{SVD}}^r, \end{aligned}$$

where $(\mathbf{V}^r)^\perp$ is the orthogonal complement of \mathbf{V}^r in \mathbb{R}^p and $\mathbf{D} := \text{diag}[\sigma_{r+1}, \dots, \sigma_p]$. The second equality clearly shows that the regularization scheme adds infinity to all singular values that correspond to data-uninformed modes. The last equality proves that infinite regularization on data-uninformed parameter subspace is the same as the truncated SVD approach.

The beauty of the truncated SVD approach is that it avoids putting any regularization on data-informed parameter directions, and hence avoids polluting inverse solutions in these directions, while annihilating data-uninformed directions. However, it is often the case that there is no clear-cut (i.e. $\sigma_k = 0$ for $k \geq r+1$) between the data-informed and data-uninformed ones, but gradual (sometimes exponential) decay of singular values of \mathbf{A} . In that case, completely removing less data-informed directions may not be ideal, as they may still contain valuable parameter directions seen by the data. Then we may want to impose finite regularization in less data-informed directions, i.e.,

$$(2.5) \quad \min \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2 + \frac{1}{2} \|\mathbf{L}(\mathbf{x} - \mathbf{x}_0)\|^2,$$

where

$$L^T L := \alpha \left(\mathbf{I} - \mathbf{V}^r (\mathbf{V}^r)^T \right) = \alpha (\mathbf{V}^r)^\perp \left((\mathbf{V}^r)^\perp \right)^T = \left[\begin{array}{c|c} 0 & 0 \\ \hline 0 & \alpha \mathbf{I} \end{array} \right] \left[\mathbf{V}^r, (\mathbf{V}^r)^\perp \right]^T.$$

Let us call this approach the **data-informed (DI) regularization method**. The inverse solution in this case reads

$$\begin{aligned} \mathbf{x}_{DI} &= \left\{ \mathbf{A}^T \mathbf{A} + \alpha \left(\mathbf{I} - \mathbf{V}^r (\mathbf{V}^r)^T \right) \right\}^{-1} \mathbf{A}^T \mathbf{y} \\ &= \left\{ \left[\mathbf{V}^r, (\mathbf{V}^r)^\perp \right] \left(\left[\begin{array}{c|c} \Sigma_r^2 & 0 \\ \hline 0 & \mathbf{D} \end{array} \right] + \left[\begin{array}{c|c} 0 & 0 \\ \hline 0 & \alpha \mathbf{I} \end{array} \right] \right) \left[\mathbf{V}^r, (\mathbf{V}^r)^\perp \right]^T \right\}^{-1} \mathbf{A}^T \mathbf{y} \\ &= \left\{ \left[\mathbf{V}^r, (\mathbf{V}^r)^\perp \right] \left(\left[\begin{array}{c|c} \Sigma_r^2 & 0 \\ \hline 0 & \mathbf{D} \end{array} \right] + \alpha \left[\begin{array}{c|c} \mathbf{I} & 0 \\ \hline 0 & \mathbf{I} \end{array} \right] - \left[\begin{array}{c|c} \alpha \mathbf{I} & 0 \\ \hline 0 & 0 \end{array} \right] \right) \left[\mathbf{V}^r, (\mathbf{V}^r)^\perp \right]^T \right\}^{-1} \mathbf{A}^T \mathbf{y}. \end{aligned}$$

The last equality suggests that the DI approach is equivalent to *first applying the same¹ (finite) regularization for all parameter directions and then removing regularizations in the data-informed directions*. A few observations are in order: 1) When

¹Note that α need not be the same for all directions.

$r = 0$, DI becomes the standard Tikhonov regularization; 2) When $\alpha \rightarrow \infty$ DI approaches the truncated SVD; and 3) when $\alpha \ll \sigma_i$ for $i \leq r$, i.e. under-regularization, Tikhonov is close to DI as the contribution of the regularization to data-informed modes is small. These observations are clearly demonstrated in Figure 3 for a 1D deconvolution with $\lambda = 0.05$ with various combination of regularization parameter α and the number of retained data-informed modes r . An important feature of the DI technique that can be seen from this result is that for each r the DI solution is robust with the regularization parameter, that is, the solution does not alter significantly, especially for moderate-to-large regularization, while Tikhonov solution is damped out as the regularization parameter increases. The last column of Figure 3 shows that for $r = 20$ the DI solution retains high frequency modes which are not regularized and thus oscillatory.

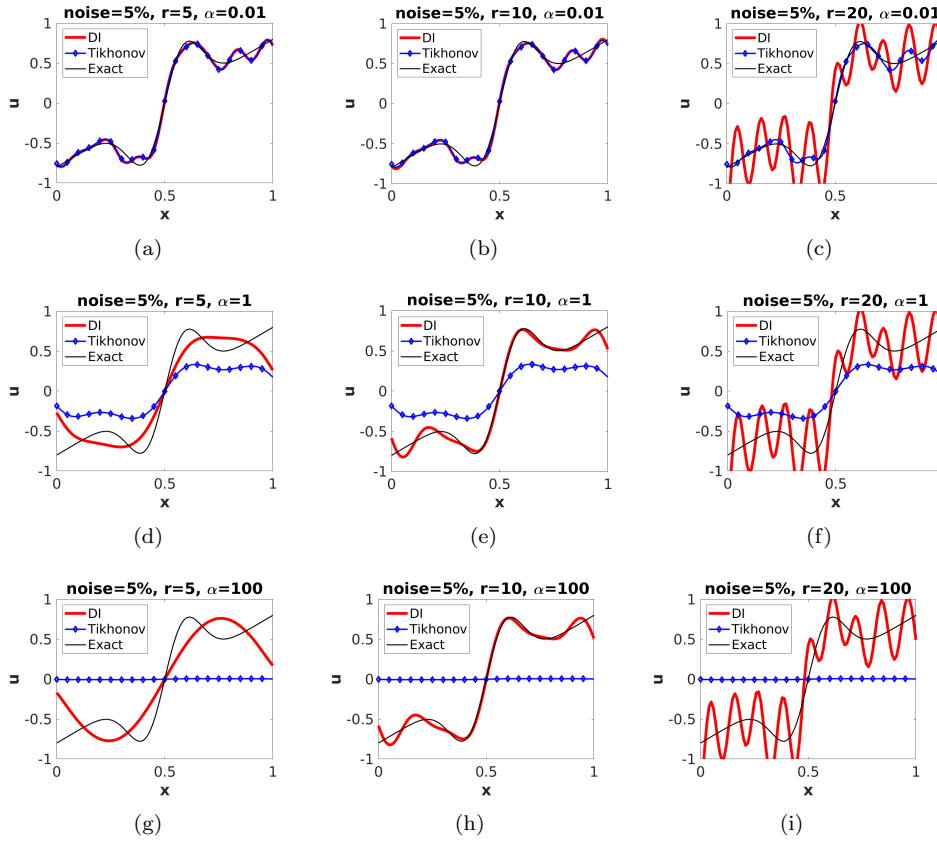


FIG. 3. Deconvolution with noise level $\lambda = 5\%$ using DI and Tikhonov regularization for various values of regularization parameter and r .

In order to gain more insights into the behavior of DI regularization, we compute the relative error between the solutions using the DI approach and the truth for a wide range of regularization parameters and a few values of r . The result is shown in Figure 4. As can be seen, when $r = 1$, DI is essentially Tikhonov, which is not surprising as all modes in the DI solution are regularized exactly the same as Tikhonov except for the first one (lowest frequency). For $r = \{5, 10\}$, the DI solution behaves the same as Tikhonov for the under-regularization regime ($\alpha < 0.01$) as expected, and

it outperforms Tikhonov for $\alpha > 0.01$ as the retained data-informed modes, which determine the quality of the deconvolution solution, are left untouched. For $r = 20$, the retained modes now also include high frequency modes and hence the DI approach is not as accurate as Tikhonov for $\alpha < 1$. For all cases with significant number of modes retained, i.e. $r > 5$, the DI solution quality is insensitive to a large range of the regularization parameter.

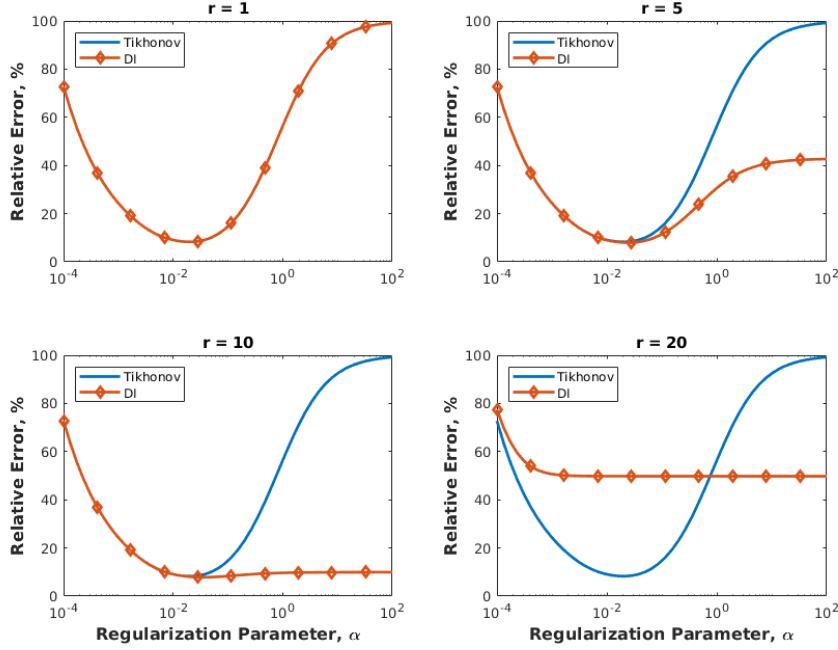


FIG. 4. Deconvolution with noise level $\lambda = 5\%$ using DI and Tikhonov regularizations for $\alpha = [10^{-4}, 10^2]$ and $r = \{1, 5, 10, 20\}$.

It may appear that we have to first construct the rank- r data-informed subspace \mathbf{V}^r before applying regularization in its orthogonal complement. However, since $\mathbf{V}^r (\mathbf{V}^r)^T$ is the orthogonal projection into the row space of \mathbf{A} when $\sigma_i = 0$ for $i > r$, i.e. $\mathbf{V}^r (\mathbf{V}^r)^T = \mathbf{A}^T (\mathbf{A} \mathbf{A}^T)^\dagger \mathbf{A}$, we can rewrite the inverse (optimization) problem (2.5) as

$$(2.6) \quad \min_{\mathbf{x}} J := \frac{1}{2} \|\mathbf{A} \mathbf{x} - \mathbf{y}\|^2 + \frac{1}{2} \|\mathbf{L}(\mathbf{x} - \mathbf{x}_0)\|^2,$$

where

$$\mathbf{L}^T \mathbf{L} := \alpha \left(\mathbf{I} - \mathbf{A}^T (\mathbf{A} \mathbf{A}^T)^\dagger \mathbf{A} \right).$$

In this form, the DI regularization approach (2.6) not only avoids using \mathbf{V}^r explicitly but also brings us to a statistical data-informed inverse framework in the next section.

2.2. A statistical data-informed (DI) inverse framework. The cost function in (2.6) can be rewritten as

$$\exp(-J) = \frac{\exp\left(-\frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2\right) \times \exp\left(-\frac{\alpha}{2} \|\mathbf{x} - \mathbf{x}_0\|^2\right)}{\exp\left(-\frac{\alpha}{2} (\mathbf{A}\mathbf{x} - \mathbf{A}\mathbf{x}_0)^T (\mathbf{A}\mathbf{A}^T)^\dagger (\mathbf{A}\mathbf{x} - \mathbf{A}\mathbf{x}_0)\right)}.$$

From a Bayesian inverse perspective [6, 13, 3, 8, 7, 12, 10], the numerator is the product of the likelihood

$$\pi_{\text{like}}(\mathbf{y}|\mathbf{x}) \propto \exp\left(-\frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2\right)$$

from the observational model (2.1) with the noise $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and the Gaussian prior

$$(2.7) \quad \pi_{\text{prior}}(\mathbf{x}) \propto \exp\left(-\frac{\alpha}{2} \|\mathbf{x} - \mathbf{x}_0\|^2\right)$$

with mean \mathbf{x}_0 and \mathbf{I}/α covariance matrix. In other words, the numerator is a Bayesian posterior with the aforementioned likelihood and Gaussian prior. *The key difference compared to the Bayesian approach is the denominator.*

We now show that the denominator is nothing more than the pushforward of the prior (2.7) via the forward map \mathbf{A} . Indeed, let $\tilde{\mathbf{y}} := \mathbf{A}\mathbf{x}$ be a random variable induced by the forward map \mathbf{A} . With $\mathbf{x} \sim \mathcal{N}(\mathbf{x}_0, \mathbf{I}/\alpha)$, $\tilde{\mathbf{y}}$ is also a Gaussian with mean $\tilde{\mathbf{y}}_0$ and covariance matrix \mathbb{C} where

$$\begin{aligned} \tilde{\mathbf{y}}_0 &:= \mathbb{E}_{\mathbf{x}}[\mathbf{A}\mathbf{x}] = \mathbf{A}\mathbf{x}_0 \\ \mathbb{C} &:= \mathbb{E}_{\mathbf{x}}[(\tilde{\mathbf{y}} - \tilde{\mathbf{y}}_0)(\tilde{\mathbf{y}} - \tilde{\mathbf{y}}_0)^T] = \mathbb{E}_{\mathbf{x}}[\mathbf{A}(\mathbf{x} - \mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0)^T \mathbf{A}^T] = \frac{1}{\alpha} \mathbf{A}\mathbf{A}^T. \end{aligned}$$

Note that it is necessary to use the pseudo-inverse for the inverse of the covariance \mathbb{C} , i.e., $\mathbb{C}^{-1} := \alpha (\mathbf{A}\mathbf{A}^T)^\dagger$ since \mathbf{A} may not have full row rank and thus the pushforward distribution can be a degenerate Gaussian.

Remark 2.1. The push-forward of the prior through the parameter-to-observable map \mathbf{A} depends on \mathbf{x} . *It is through this push-forward term that the data-informed (DI) approach learns the data-informed parameter directions.* Indeed, this new approach, through the push-forward term, changes the original prior

$$\exp\left(-\frac{\alpha}{2} (\mathbf{x} - \mathbf{x}_0)^T \mathbf{I} (\mathbf{x} - \mathbf{x}_0)\right)$$

to the new one

$$\exp\left(-\frac{\alpha}{2} (\mathbf{x} - \mathbf{x}_0)^T \left[\mathbf{I} - \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^\dagger \mathbf{A}\right] (\mathbf{x} - \mathbf{x}_0)\right)$$

in such a way that the new prior leaves the data-informed directions, i.e. the row space of \mathbf{A} , untouched, and hence only regularizes data-uninformed directions. The data-informed approach accomplishes this by the pushforward of the prior via the parameter-to-observable map \mathbf{A} .

We can now define the DI posterior as

$$(2.8) \quad \pi_{\text{DI}}(\mathbf{x}|\mathbf{y}) = \frac{\pi_{\text{like}}(\mathbf{y}|\mathbf{x}) \times \pi_{\text{prior}}(\mathbf{x})}{\mathbf{A}_{\#} \pi_{\text{prior}}(\mathbf{x})},$$

where $\mathbf{A}_{\#} \pi_{\text{prior}}(\mathbf{x})$ denotes the pushforward of $\pi_{\text{prior}}(\mathbf{x})$ via the parameter-to-observable map \mathbf{A} .

We have constructively derived the DI approach by modifying the truncated SVD method and Gaussian prior with scaled-identity covariance matrix. In practice, the prior can be more informative about the correlations among components of \mathbf{x} and in that case the covariance matrix is no longer an identity matrix. Let us denote by $\pi_{\text{prior}}(\mathbf{x}) = \mathcal{N}(\mathbf{x}_0, \mathbf{\Gamma}/\alpha)$ the Gaussian prior with covariance matrix $\mathbf{\Gamma}/\alpha$. Let us also consider a more general data distribution where, for a given parameter \mathbf{x} , the data is distributed by the Gaussian $\mathcal{N}(\mathbf{Ax}, \mathbf{\Lambda})$. In order to use most of the above results, let us whiten both the parameter and observations. In particular, $\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{y}$ is the whitened observations (inducing $\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{A}$ as the new parameter-to-observable forward map) and $\mathbf{\Gamma}^{-\frac{1}{2}}\mathbf{x}$ is the prior-whitened parameter. The pushforward of the prior via $\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{A}$ now reads²

$$(2.9) \quad \mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{A}_{\#} \pi_{\text{prior}}(\mathbf{x}) = \mathcal{N}\left(\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{Ax}_0, \frac{1}{\alpha}\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T\mathbf{\Lambda}^{-\frac{1}{2}}\right),$$

The DI posterior (2.8) with whitened parameter, whitened observations, and induced parameter-to-observable map now becomes

$$(2.10) \quad \pi_{\text{DI}}(\mathbf{x}|\mathbf{y}) = \frac{\pi_{\text{like}}(\mathbf{y}|\mathbf{x}) \times \pi_{\text{prior}}(\mathbf{x})}{\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{A}_{\#} \pi_{\text{prior}}(\mathbf{x})} \\ \propto \frac{\exp\left(-\frac{1}{2}\left\|\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{Ax} - \mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{y}\right\|^2\right) \times \exp\left(-\frac{\alpha}{2}\left\|\mathbf{\Gamma}^{-\frac{1}{2}}\mathbf{x} - \mathbf{\Gamma}^{-\frac{1}{2}}\mathbf{x}_0\right\|^2\right)}{\exp\left(-\frac{\alpha}{2}\left\|\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{Ax} - \mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{Ax}_0\right\|^2_{\left(\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T\mathbf{\Lambda}^{-\frac{1}{2}}\right)^{\dagger}}\right)},$$

which, after writing the push-forward measure in terms of the whitened parameter, reads

$$\pi_{\text{DI}}(\mathbf{x}|\mathbf{y}) \propto \frac{\exp\left(-\frac{1}{2}\left\|\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{Ax} - \mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{y}\right\|^2\right) \times \exp\left(-\frac{\alpha}{2}\left\|\mathbf{\Gamma}^{-\frac{1}{2}}\mathbf{x} - \mathbf{\Gamma}^{-\frac{1}{2}}\mathbf{x}_0\right\|^2\right)}{\exp\left(-\frac{\alpha}{2}\left\|\mathbf{\Gamma}^{-\frac{1}{2}}\mathbf{x} - \mathbf{\Gamma}^{-\frac{1}{2}}\mathbf{x}_0\right\|^2_{\mathbf{\Gamma}^{\frac{1}{2}}\mathbf{A}^T\mathbf{\Lambda}^{\frac{1}{2}}\left(\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{A}\mathbf{\Gamma}\mathbf{A}^T\mathbf{\Lambda}^{-\frac{1}{2}}\right)^{\dagger}\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{A}\mathbf{\Gamma}^{\frac{1}{2}}}\right)},$$

or equivalently

$$(2.11) \quad -\log(\pi_{\text{DI}}(\mathbf{x}|\mathbf{y})) \propto \frac{1}{2}\left\|\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{Ax} - \mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{y}\right\|^2 + \frac{1}{2}\left\|L\left(\mathbf{\Gamma}^{-\frac{1}{2}}\mathbf{x} - \mathbf{\Gamma}^{-\frac{1}{2}}\mathbf{x}_0\right)\right\|^2,$$

²Note that using the modified forward map $\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{A}$ is not necessary as using the original map \mathbf{A} yields the same result.

where

$$\begin{aligned}
 L^T L &= \alpha \left(\mathbf{I} - \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{A}^T \mathbf{\Lambda}^{-\frac{1}{2}} \left(\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma} \mathbf{A}^T \mathbf{\Lambda}^{-\frac{1}{2}} \right)^\dagger \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}} \right) \\
 &= \alpha \left(\mathbf{I} - \mathbf{V}^n \mathbf{V}^{nT} \right) = \alpha (\mathbf{V}^n)^\perp \left((\mathbf{V}^n)^\perp \right)^T \\
 (2.12) \quad &= \left[\mathbf{V}^n, (\mathbf{V}^n)^\perp \right] \left[\begin{array}{c|c} 0 & 0 \\ \hline 0 & \alpha \mathbf{I} \end{array} \right] \left[\mathbf{V}^n, (\mathbf{V}^n)^\perp \right]^T,
 \end{aligned}$$

where \mathbf{V}^n contains the first n right singular vectors of the following SVD

$$(2.13) \quad \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}} := \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T.$$

As can be seen, the push-forward measure seeks to find the first n columns of \mathbf{V} associated with the n non-zero singular values. The DI method then avoids regularizing these “data-informed directions” \mathbf{V}^n . In other words, in the whitened parameter, the induced regularization by the prior is identity and the DI approach removes regularization in the parameter subspace spanned by \mathbf{V}^n . From (2.13) it is clear that \mathbf{V}^n now depends on both the prior covariance $\mathbf{\Gamma}$ and the observational covariance $\mathbf{\Lambda}$ in addition to \mathbf{A} . *So how do we understand the parameter subspace spanned by \mathbf{V}^n , and hence the DI approach?* To that end, let us define $\mathbf{\Sigma}^{\frac{1}{2}}$ to be the same as $\mathbf{\Sigma}$ except on the main the diagonal where $\mathbf{\Sigma}^{\frac{1}{2}}(i, i) = \sqrt{\mathbf{\Sigma}(i, i)} = \sqrt{\sigma_i}$ (note that $\mathbf{\Sigma}^{\frac{1}{2}}$ is nothing more than the square-root of $\mathbf{\Sigma}$ when $\mathbf{\Sigma}$ is a square matrix). Let $\mathbf{\Psi}$ be the first n rows of $\mathbf{\Sigma}^{\frac{1}{2}}$ and $\mathbf{\Phi}$ be the first n columns of $\mathbf{\Sigma}^{\frac{1}{2}}$. Clearly, by definition $\mathbf{\Psi}(i, i) = \mathbf{\Phi}(i, i) = \sqrt{\sigma_i}$ for $i \leq n$.

Let us define the following maps

$$(2.14) \quad \mathbf{z} := \mathbf{T} \mathbf{x}, \quad \text{where} \quad \mathbf{T} := \mathbf{\Psi} \mathbf{V}^T \mathbf{\Gamma}^{-\frac{1}{2}},$$

$$(2.15) \quad \mathbf{w} := \mathbf{S} \mathbf{y}, \quad \text{where} \quad \mathbf{S} := \mathbf{\Phi}^T \mathbf{U}^T \mathbf{\Lambda}^{-\frac{1}{2}}.$$

where \mathbf{z} are the first n coordinates of \mathbf{x} in \mathbf{V} , after whitening via $\mathbf{\Gamma}^{-\frac{1}{2}}$ and then being scaled by $\mathbf{\Psi}$. Similarly, \mathbf{w} are the first n coordinates of \mathbf{y} in \mathbf{U} , after whitening via $\mathbf{\Lambda}^{-\frac{1}{2}}$ and then being scaled by $\mathbf{\Phi}$. The map \mathbf{T} pushes forward the prior in \mathbf{x} to the prior in \mathbf{z} as

$$(2.16) \quad \pi_{\text{prior}}(\mathbf{z}) \sim \exp \left(-\frac{1}{2} \sum_{i=1}^n \sigma_i^{-1} (z_i - \bar{z}_i)^2 \right),$$

where $\bar{\mathbf{z}} = \mathbf{T} \mathbf{x}_0$. Similarly, given \mathbf{x} (and hence \mathbf{z}), the induced likelihood in terms of \mathbf{w} is given by

$$(2.17) \quad \pi_{\text{like}}(\mathbf{w}|\mathbf{z}) \sim \exp \left(-\frac{1}{2} \sum_{i=1}^n \sigma_i^{-1} (\mathbf{w}_i - \sigma_i \mathbf{z}_i)^2 \right).$$

As can be seen from (2.16) and (2.17), the maps \mathbf{T} and \mathbf{S} transform the original parameter \mathbf{x} and original data \mathbf{y} to new parameter \mathbf{z} and new data \mathbf{w} . Two observations are in order: 1) though in general the original parameter and data dimensions are different, the new parameter and data have the same dimension; and 2) the new data \mathbf{w} and new parameter \mathbf{z} , up to the difference in the mean, have the

same distribution. In particular, both \mathbf{z} and \mathbf{w} are \mathbb{R}^n -vector of independent Gaussian distributions with diagonal covariance matrix $\Theta \in \mathbb{R}^{n \times n}$ with $\Theta_{ii} = \sigma_i$. Both \mathbf{z}_i and \mathbf{w}_i , up to the difference in the mean, are the same Gaussian distribution with variance σ_i . Since $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n > 0$, the independent random variable \mathbf{z}_i (and hence \mathbf{w}_i) are ranked from the one with most variance to the one with least variance.

Let us call the i th column of \mathbf{U} , namely \mathbf{U}_i , the i th important direction in the data space, and the i th column of \mathbf{V} , namely \mathbf{V}_i , the i th important direction in the parameter space. Let us also rank the degree of importance of \mathbf{U}_i and \mathbf{V}_i by the magnitude of σ_i . It follows that the transformations \mathbf{T} and \mathbf{S} map the original parameter \mathbf{x} and data \mathbf{y} into new parameter \mathbf{z} and data \mathbf{w} in which the corresponding parameter \mathbf{z}_i and data \mathbf{w}_i are equally important. This is similar to the concept of balanced transformation in control theory (see, e.g., [4, 1] and the references therein). The new parameter \mathbf{z} is thus equally data-informed and prior-informed. In particular \mathbf{z}_i is equally less data-informed and prior-informed relatively to \mathbf{z}_j for $j < i$.

The DI method thus regularizes only the (equally) data-uninformed and prior-uninformed parameters/directions.

2.3. Properties of the DI regularization approach.

2.3.1. Deterministic properties. It is easy to see the optimality condition of the optimization problem $\max_{\mathbf{x}} \log(\pi_{\text{DI}}(\mathbf{x}|\mathbf{y}))$ is given by

$$(2.18) \quad \mathbf{H}\mathbf{x}_{\text{DI}} = \mathbf{b},$$

where

$$\begin{aligned} \mathbf{H} &:= \left\{ \mathbf{A}^T \Lambda^{-1} \mathbf{A} + \alpha \left[\Gamma^{-1} - \mathbf{A}^T \Lambda^{-\frac{1}{2}} \left(\Lambda^{-\frac{1}{2}} \mathbf{A} \Gamma \mathbf{A}^T \Lambda^{-\frac{1}{2}} \right)^\dagger \Lambda^{-\frac{1}{2}} \mathbf{A} \right] \right\}, \\ \mathbf{b} &:= \mathbf{A}^T \Lambda^{-1} \mathbf{y} + \alpha \left[\Gamma^{-1} - \mathbf{A}^T \Lambda^{-\frac{1}{2}} \left(\Lambda^{-\frac{1}{2}} \mathbf{A} \Gamma \mathbf{A}^T \Lambda^{-\frac{1}{2}} \right)^\dagger \Lambda^{-\frac{1}{2}} \mathbf{A} \right] \mathbf{x}_0 \end{aligned}$$

In order to solve the optimality condition (2.18) in practice, we can use the rank- r approximation

$$(2.19) \quad \Lambda^{-\frac{1}{2}} \mathbf{A} \Gamma^{\frac{1}{2}} = \mathbf{U}^n \Sigma^n (\mathbf{V}^n)^T \approx \mathbf{U}^r \Sigma^r (\mathbf{V}^r)^T$$

for the push-forward matrix $\mathbf{A}^T \Lambda^{-\frac{1}{2}} \left(\Lambda^{-\frac{1}{2}} \mathbf{A} \Gamma \mathbf{A}^T \Lambda^{-\frac{1}{2}} \right)^\dagger \Lambda^{-\frac{1}{2}} \mathbf{A}$, where n is largest index for which $\sigma_n > 0$. Thus rank- r approximations for \mathbf{H} and \mathbf{y} are given by

$$\begin{aligned} \mathbf{H}^r &:= \mathbf{A}^T \Lambda^{-1} \mathbf{A} + \alpha \left(\Gamma^{-1} - \Gamma^{-\frac{1}{2}} \mathbf{V}^r (\mathbf{V}^r)^T \Gamma^{-\frac{1}{2}} \right), \\ \mathbf{b}^r &:= \Gamma^{-\frac{1}{2}} \mathbf{V}^n \Sigma^n (\mathbf{U}^n)^T \Lambda^{-\frac{1}{2}} \mathbf{y} + \alpha \left[\Gamma^{-1} - \Gamma^{-\frac{1}{2}} \mathbf{V}^r (\mathbf{V}^r)^T \Gamma^{-\frac{1}{2}} \right] \mathbf{x}_0. \end{aligned}$$

Note that we don't perform low rank approximation for the term $\mathbf{A}^T \Lambda^{-1} \mathbf{y}$ in \mathbf{y} as it requires only a matrix-vector product. We also leave the first term in \mathbf{H}^r as is, since we invert \mathbf{H}^r using the conjugate gradient (CG) method which requires only matrix-vector products. In the numerical results section, section 3, we present a nested optimization method that avoids the low rank approximation altogether. The analysis of such method is, however, more technical and thus left for future work. The rank- r approximation to the solution of the optimality condition (2.18) is defined as

$$(2.20) \quad \mathbf{H}^r \mathbf{x}_{\text{DI}}^r = \mathbf{b}^r,$$

which is solved approximately using the conjugate gradient (CG) method (see Section 3). For the following analysis, however, we compute the inverse of \mathbf{H}^r analytically. To that end, we rewrite \mathbf{H}^r in terms of n singular vectors corresponding to the n nonzero singular values as

$$\mathbf{H}^r = \alpha \mathbf{\Gamma}^{-\frac{1}{2}} \left[\mathbf{I} + \mathbf{V}^n \mathbf{D}^n (\mathbf{V}^n)^T \right] \mathbf{\Gamma}^{-\frac{1}{2}},$$

where \mathbf{D}^n is an $n \times n$ diagonal matrix with $\mathbf{D}^n(i, i) = (\sigma_i^2 - \alpha) / \alpha$ for $i \leq r$ and $\mathbf{D}^n(i, i) = \sigma_i^2 / \alpha$ for $r < i \leq n$.

LEMMA 2.2. *The DI solution with r data-informed modes reads*

$$(2.21) \quad \mathbf{x}_{DI}^r := \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{V}^n \mathbf{\Theta}^n (\mathbf{U}^n)^T \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{y} + \left[\mathbf{I} - \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{V}^n \bar{\mathbf{T}}^n (\mathbf{V}^n)^T \mathbf{\Gamma}^{-\frac{1}{2}} \right] \mathbf{x}_0,$$

where $\mathbf{\Theta}^n$ is an $n \times n$ diagonal matrix with $\mathbf{\Theta}^n(i, i) = \sigma_i^{-1}$ for $i \leq r$ and $\mathbf{\Theta}^n(i, i) = \sigma_i / (\sigma_i^2 + \alpha)$ for $r < i \leq n$. Here, $\bar{\mathbf{T}}^n$ is an $n \times n$ diagonal matrix with $\bar{\mathbf{T}}^n(i, i) = 1$ for $i \leq r$ and $\bar{\mathbf{T}}^n(i, i) = \sigma_i^2 / (\sigma_i^2 + \alpha)$ for $r < i \leq n$. Furthermore,

$$(2.22) \quad \mathbf{A} \mathbf{x}_{DI}^r = \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{U}^n (\mathbf{U}^n)^T \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{y}.$$

Proof. Using a Woodbury formula we have

$$(2.23) \quad (\mathbf{H}^r)^{-1} = \frac{1}{\alpha} \mathbf{\Gamma}^{\frac{1}{2}} \left[\mathbf{I} - \mathbf{V}^n \bar{\mathbf{D}}_{DI}^{n,r} (\mathbf{V}^n)^T \right] \mathbf{\Gamma}^{\frac{1}{2}},$$

where $\bar{\mathbf{D}}_{DI}^{n,r}$ is an $n \times n$ diagonal matrix with $\bar{\mathbf{D}}_{DI}^{n,r}(i, i) = (\sigma_i^2 - \alpha) / \sigma_i^2$ for $i \leq r$ and $\bar{\mathbf{D}}_{DI}^{n,r}(i, i) = \sigma_i^2 / (\sigma_i^2 + \alpha)$ for $r < i \leq n$. The computation of the product $(\mathbf{H}^r)^{-1} \mathbf{y}^r$ to arrive at the assertion is straightforward algebra manipulation and hence omitted. \square

The result (2.22) shows that the image of the DI solution \mathbf{x}_{DI} through the parameter-to-observable map is exactly the data if $\mathbf{U}^n (\mathbf{U}^n)^T = \mathbf{I}$ or $\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{y}$ resides in the column space of \mathbf{U}^n . This happens, for example, when \mathbf{A} has full row-rank and the number of data is not more than the dimension of the parameter, i.e., $d \leq p$. In this case, retaining all modes corresponding to non-zero singular values in the DI solution makes the data misfit vanish, that is, the DI solution in this case would match the noise, which is undesirable. As discussed in Section 2.1, r should be smaller than n for the solution to be meaningful. Let us define

$$r_\varepsilon := \max \{i : 1 \leq i \leq n \text{ and } \sigma_i \geq \varepsilon\},$$

for some $\varepsilon > 0$, and

$$\mathcal{R}_\varepsilon := (\mathbf{H}^{r_\varepsilon})^{-1} \mathbf{A}^T \mathbf{\Lambda}^{-\frac{1}{2}}.$$

THEOREM 2.3. *For any $\varepsilon > 0$ and $\alpha > 0$, consider the inverse problem*

$$(2.24) \quad \min_{\mathbf{x}} \mathcal{J} = \frac{1}{2} \left\| \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{x} - \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{y} \right\|^2 + \frac{1}{2} \left\| L \mathbf{\Gamma}^{-\frac{1}{2}} (\mathbf{x} - \mathbf{x}_0) \right\|^2,$$

using the DI technique with rank- r_ε approximation, where

$$L^T L = \alpha \left(\mathbf{I} - \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{A}^T \mathbf{\Lambda}^{-\frac{1}{2}} \left(\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma} \mathbf{A}^T \mathbf{\Lambda}^{-\frac{1}{2}} \right)^\dagger \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}} \right).$$

The following hold:

- i) The inverse problem with rank r_ε DI technique is well-posed.
 ii) Suppose that the nullspace of \mathbf{A} is trivial, i.e. $\mathcal{N}(\mathbf{A}) = \{0\}$, then the DI technique is a regularization strategy in the following sense

$$\lim_{\varepsilon \rightarrow 0} \mathcal{R}_\varepsilon \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{x} = \mathbf{x}.$$

- iii) If $\alpha = \mathcal{O}(\varepsilon)$ and $\mathcal{N}(\mathbf{A}) = \{0\}$, then the rank- r_ε DI technique is an admissible regularization method.

Proof. From Lemma 2.2 we see that the DI solution $\mathbf{x}_{\text{DI}}^{r_\varepsilon}$ is unique and furthermore

$$\|\mathbf{x}_{\text{DI}}^{r_\varepsilon}\| \leq \beta(\varepsilon, \alpha) \left\| \mathbf{\Gamma}^{\frac{1}{2}} \right\| \left\| \mathbf{\Lambda}^{-\frac{1}{2}} \right\| \|\mathbf{y}\| + \left(1 + \sqrt{\kappa(\mathbf{\Gamma})}\right) \|\mathbf{x}_0\|,$$

where $\kappa(\mathbf{\Gamma})$ denotes the condition number of $\mathbf{\Gamma}$, $\beta(\varepsilon, \alpha)$ is a constant defined as

$$\beta(\varepsilon, \alpha) := \frac{1}{\min_{r_\varepsilon < i \leq n} \{r_\varepsilon, \sigma_i + \alpha/\sigma_i\}},$$

which shows the the DI solution is stable, which in turn proves i). To see assertion ii), we use the definition of \mathcal{R}_ε and the SVD of $\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}}$ to arrive at

$$\begin{aligned} \mathcal{R}_\varepsilon \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} &= \frac{1}{\alpha} \mathbf{\Gamma}^{\frac{1}{2}} \left[\mathbf{I} - \mathbf{V}^n \overline{\mathbf{D}}^{n, r_\varepsilon} (\mathbf{V}^n)^T \right] \mathbf{V}^n (\mathbf{\Sigma}^n)^2 (\mathbf{V}^n)^T \mathbf{\Gamma}^{-\frac{1}{2}} = \\ &\quad \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{V}^n \left[\begin{array}{c|c} \mathbf{I} & 0 \\ 0 & \text{diag} \left(\frac{\sigma_i^2}{\sigma_i^2 + \alpha} \right)_{r_\varepsilon < i \leq n} \end{array} \right] (\mathbf{V}^n)^T \mathbf{\Gamma}^{-\frac{1}{2}}, \end{aligned}$$

which implies

$$\lim_{\varepsilon \rightarrow 0} \mathcal{R}_\varepsilon \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{x} = \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{V}^n \mathbf{I} (\mathbf{V}^n)^T \mathbf{\Gamma}^{-\frac{1}{2}} \mathbf{x} = \mathbf{x},$$

where we have used the fact that $r_\varepsilon \rightarrow n$ as $\varepsilon \rightarrow 0$, and that $\mathbf{V}^n (\mathbf{V}^n)^T = \mathbf{I}$ since $\mathcal{N}(\mathbf{A}) = \{0\}$.

For assertion iii), it is sufficient to show that

$$\sup_{\mathbf{y}} \left\{ \left\| \mathcal{R}_\varepsilon \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{y} - \mathbf{x} \right\| : \left\| \mathbf{\Lambda}^{-\frac{1}{2}} (\mathbf{A} \mathbf{x} - \mathbf{y}) \right\| \leq \varepsilon \right\} \rightarrow 0 \text{ as } \varepsilon \rightarrow 0,$$

for any \mathbf{x} . We have

$$\begin{aligned} \left\| \mathcal{R}_\varepsilon \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{y} - \mathbf{x} \right\| &\leq \left\| \mathcal{R}_\varepsilon \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{x} - \mathbf{x} \right\| + \left\| \mathcal{R}_\varepsilon \mathbf{\Lambda}^{-\frac{1}{2}} (\mathbf{A} \mathbf{x} - \mathbf{y}) \right\| \\ &\leq \left\| \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{V}^n \left[\begin{array}{c|c} 0 & 0 \\ 0 & \text{diag} \left(\frac{-\alpha}{\sigma_i^2 + \alpha} \right)_{r_\varepsilon < i \leq n} \end{array} \right] (\mathbf{V}^n)^T \mathbf{\Gamma}^{-\frac{1}{2}} \right\| \|\mathbf{x}\| + \|\mathcal{R}_\varepsilon\| \varepsilon \\ &\leq \frac{\alpha}{\sigma_n^2 + \alpha} \sqrt{\kappa(\mathbf{\Gamma})} \|\mathbf{x}\| + \varepsilon \left\| \mathbf{\Gamma}^{\frac{1}{2}} \right\| \left\| \left[\begin{array}{c|c} \text{diag} \left(\frac{1}{\sigma_i} \right)_{i < r_\varepsilon} & 0 \\ 0 & \text{diag} \left(\frac{\sigma_i}{\sigma_i^2 + \alpha} \right)_{r_\varepsilon < i \leq n} \end{array} \right] \right\| \\ &\leq \frac{\alpha}{\sigma_n^2 + \alpha} \sqrt{\kappa(\mathbf{\Gamma})} \|\mathbf{x}\| + \varepsilon \sigma_n^{-1} \left\| \mathbf{\Gamma}^{\frac{1}{2}} \right\|, \end{aligned}$$

where we have used the result from ii), definition of \mathcal{R}_ε , and the orthonormality of \mathbf{V} and \mathbf{U} . Using the assumption $\alpha = \mathcal{O}(\varepsilon)$ concludes the proof. \square

Remark 2.4. Note that most of the above arguments are still valid for infinite dimensional setting, i.e. $p = \infty$, assuming that $\mathbf{\Gamma}$ is a trace class. Indeed, $\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}}$ is then a compact operator and we can invoke the infinite dimensional singular value decomposition [2] for $\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}}$. Note that all the matrices are now interpreted as operators, transpose operator (superscript T) as adjoint operator, and $\mathbf{\Gamma}^{-\frac{1}{2}}$ as pseudo-inverse if $\mathcal{N}(\mathbf{\Gamma}) \neq \{0\}$. We leave out the details for the sake of brevity.

2.3.2. Statistical properties. Now we discuss some probabilistic aspects of the DI prior and the DI posterior. Since the regularization parameter α plays no role in the following discussion, we absorb it into $\mathbf{\Gamma}$. We define the DI prior as

$$(2.25) \quad \pi_{\text{DI-prior}}(\mathbf{x}) \sim \exp \left\{ -\frac{1}{2} \left\| L \mathbf{\Gamma}^{-\frac{1}{2}} (\mathbf{x} - \mathbf{x}_0) \right\|^2 \right\}.$$

From (2.12), the DI prior (pseudo-) inverse covariance is given by

$$\begin{aligned} (\mathbb{C}^n)^\dagger &:= \mathbf{\Gamma}^{-\frac{1}{2}} \left[\mathbf{I} - \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{A}^T \mathbf{\Lambda}^{-\frac{1}{2}} \left(\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma} \mathbf{A}^T \mathbf{\Lambda}^{-\frac{1}{2}} \right)^\dagger \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}} \right] \mathbf{\Gamma}^{-\frac{1}{2}} \\ &= \mathbf{\Gamma}^{-\frac{1}{2}} \left[\mathbf{I} - \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{A}^T \left(\mathbf{A} \mathbf{\Gamma} \mathbf{A}^T \right)^\dagger \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}} \right] \mathbf{\Gamma}^{-\frac{1}{2}} = \mathbf{\Gamma}^{-\frac{1}{2}} (\mathbf{V}^n)^\perp \left((\mathbf{V}^n)^\perp \right)^T \mathbf{\Gamma}^{-\frac{1}{2}}, \end{aligned}$$

where we have used the fact that $\mathbf{\Lambda}$ is invertible in the second equality. Thus $\mathbf{\Lambda}$ actually contribute to neither the DI prior nor its rank- r version

$$(\mathbb{C}^r)^\dagger := \mathbf{\Gamma}^{-\frac{1}{2}} (\mathbf{V}^r)^\perp \left((\mathbf{V}^r)^\perp \right)^T \mathbf{\Gamma}^{-\frac{1}{2}}.$$

The rank- r DI covariance thus reads

$$(2.26) \quad \mathbb{C}^r := \mathbf{\Gamma}^{\frac{1}{2}} (\mathbf{V}^r)^\perp \left[(\mathbf{V}^r)^\perp \right]^T \mathbf{\Gamma}^{\frac{1}{2}} = \mathbf{\Gamma}^{\frac{1}{2}} \left(\mathbf{I} - \mathbf{V}^r (\mathbf{V}^r)^T \right) \mathbf{\Gamma}^{\frac{1}{2}}$$

which is clearly symmetric positive semidefinite in \mathbb{R}^p , though degenerate.³ The DI-prior (2.25) is not a well defined density in \mathbb{R}^p , that is, it is not absolutely continuous with respect to the Lebesgue measure in \mathbb{R}^p . This is not surprising as we argue above the the DI-prior is the prior on the less data-informed directions. Let us define

$$\mathbf{z}^\perp := \mathbf{T}^\perp \mathbf{x}, \quad \text{where} \quad \mathbf{T}^\perp := \left((\mathbf{V}^r)^\perp \right)^T \mathbf{\Gamma}^{-\frac{1}{2}}.$$

THEOREM 2.5. *The following hold true:*

- i) \mathbf{z} and \mathbf{z}^\perp are distributed by the push-forward density of the prior through \mathbf{T} and \mathbf{T}^\perp , respectively. In particular, $\mathbf{z} \sim \mathcal{N}(\mathbf{T} \mathbf{x}_0, \mathbf{I})$ and $\mathbf{z}^\perp \sim \mathcal{N}(\mathbf{T}^\perp \mathbf{x}_0, \mathbf{I})$.
- ii) the DI-prior density is the density of \mathbf{z}^\perp , and hence is well-defined.
- iii) The DI-prior density is the conditional density of \mathbf{x} given \mathbf{z} .

Proof. Assertion i) is straightforward. To see the the second assertion, we note that the density of \mathbf{z}^\perp can be written as

$$\begin{aligned} \exp \left\{ -\frac{1}{2} \left\| \mathbf{z}^\perp - \mathbf{T}^\perp \mathbf{x}_0 \right\|^2 \right\} &= \exp \left\{ -\frac{1}{2} (\mathbf{x} - \mathbf{x}_0) \left(\mathbf{T}^\perp \right)^T \mathbf{T}^\perp (\mathbf{x} - \mathbf{x}_0) \right\} \\ &= \exp \left\{ -\frac{1}{2} (\mathbf{x} - \mathbf{x}_0) \mathbf{\Gamma}^{-\frac{1}{2}} (\mathbf{V}^r)^\perp \left((\mathbf{V}^r)^\perp \right)^T \mathbf{\Gamma}^{-\frac{1}{2}} (\mathbf{x} - \mathbf{x}_0) \right\}, \end{aligned}$$

³The nullspace of \mathbb{C}^r : $\mathcal{N}(\mathbb{C}^r) := \left\{ \mathbf{x} : \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{x} \in \mathcal{R}(\mathbf{V}^r) \right\}$, where $\mathcal{R}(\cdot)$ denotes the range space.

which is exactly the DI prior (2.25). In other words, we have shown that the DI prior is a well defined density on \mathbf{z}^\perp . To see assertion *iii*) we observe that

$$\pi_{\text{prior}}(\mathbf{x}) = \pi_{\text{prior}}\left(\mathbf{V}^r \mathbf{z} + (\mathbf{V}^r)^\perp \mathbf{z}^\perp\right),$$

and thus

$$\pi_{\text{prior}}(\mathbf{x}|\mathbf{z}) = \frac{\pi_{\text{prior}}(\mathbf{x})}{\pi(\mathbf{z})} = \exp\left\{-\frac{1}{2}(\mathbf{x} - \mathbf{x}_0) \boldsymbol{\Gamma}^{-\frac{1}{2}} (\mathbf{V}^r)^\perp \left((\mathbf{V}^r)^\perp\right)^T \boldsymbol{\Gamma}^{-\frac{1}{2}} (\mathbf{x} - \mathbf{x}_0)\right\},$$

which is exactly the DI prior since $\pi(\mathbf{z}) = \mathcal{N}(\mathbf{T}\mathbf{x}_0, \mathbf{I})$ is exactly the push-forward density of $\pi_{\text{prior}}(\mathbf{x})$ via the map \mathbf{T} . \square

Remark 2.6. Note that the above decomposition of \mathbf{x} into \mathbf{z} and \mathbf{z}^\perp , through the maps \mathbf{T} and \mathbf{T}^\perp , is still valid for infinite dimensional settings. However, \mathbf{z}^\perp would be distributed by an infinite dimensional Gaussian measure with identity covariance operator, which is not a valid Gaussian measure. A more general understanding of the DI prior is through disintegration. Indeed, under mild conditions on the map \mathbf{T} and its push-forward measure of the prior measure, the DI prior is nothing more than a disintegration of the prior measure via the map \mathbf{T} and this view is also valid for infinite dimensional settings.

To quantify the uncertainty in the DI inverse solution (2.21), we can use the covariance matrix of the DI posterior (2.10). For linear inverse problems with Gaussian prior and Gaussian noise—the problems considered in this chapter—the covariance matrix is exactly the inverse of the Hessian. For rank- r DI approach, the DI posterior covariance matrix $\mathcal{C}_{\text{DI}}^{\text{post}}$ is given in (2.23), i.e.,

$$(2.27) \quad \mathcal{C}_{\text{DI}}^{\text{post}} = \frac{1}{\alpha} \boldsymbol{\Gamma} - \frac{1}{\alpha} \boldsymbol{\Gamma}^{\frac{1}{2}} \mathbf{V}^n \overline{\mathbf{D}}_{\text{DI}}^{n,r} (\mathbf{V}^n)^T \boldsymbol{\Gamma}^{\frac{1}{2}}$$

It is easy to see that the covariance matrix corresponding to the Tikhonov regularization is given by

$$(2.28) \quad \mathcal{C}_{\text{Tik}}^{\text{post}} = \frac{1}{\alpha} \boldsymbol{\Gamma} - \frac{1}{\alpha} \boldsymbol{\Gamma}^{\frac{1}{2}} \mathbf{V}^n \overline{\mathbf{D}}_{\text{Tik}}^{n,r} (\mathbf{V}^n)^T \boldsymbol{\Gamma}^{\frac{1}{2}},$$

where both $\overline{\mathbf{D}}_{\text{DI}}^{n,r}$ and $\overline{\mathbf{D}}_{\text{Tik}}^{n,r}$ are diagonal matrices given in Table 1. Note that we have used α as the magnitude of the regularization to study the robustness and accuracy of all methods. If not needed, α can be straightforwardly absorbed into $\boldsymbol{\Gamma}$, and hence σ_i^2 ; in that case α is simply replaced by 1 everywhere (including those in Table 1) it appears. As can be seen, $\overline{\mathbf{D}}_{\text{Tik}}^{n,r}(i, i)$ is always non-negative for all i ,

	$i \leq r$	$r < i \leq n$
DI posterior	$\overline{\mathbf{D}}_{\text{DI}}^{n,r}(i, i) = \frac{\sigma_i^2 - \alpha}{\sigma_i^2}$	$\overline{\mathbf{D}}_{\text{DI}}^{n,r}(i, i) = \frac{\sigma_i^2}{\sigma_i^2 + \alpha}$
Tikhonov posterior	$\overline{\mathbf{D}}_{\text{Tik}}^{n,r}(i, i) = \frac{\sigma_i^2}{\sigma_i^2 + \alpha}$	$\overline{\mathbf{D}}_{\text{Tik}}^{n,r}(i, i) = \frac{\sigma_i^2}{\sigma_i^2 + \alpha}$

TABLE 1

The difference between the DI and the Tikhonov Covariance matrices.

while $\overline{\mathbf{D}}_{\text{DI}}^{n,r}(i, i)$ is negative when $\sigma_i^2 < \alpha$ for $i \leq r$. That is, while the Tikhonov

posterior uncertainty, $\mathcal{C}_{\text{Tik}}^{\text{post}}$ (Bayesian posterior with standard Gaussian prior), is always smaller than the prior uncertainty $\mathbf{\Gamma}$ no matter how much informed the data is, the DI posterior uncertainty could be higher than the prior counterpart if the data supports this. In other words, standard (or typical) Gaussian prior does not allow the data to increase the uncertainty, and hence prone to producing over-confident results (see Section 3). The DI prior, on the other hand, takes the parameter-to-observable map (the proxy to the data) into account, and thus along parameter directions that are more data-informed, i.e. $\sigma_i^2 \geq \alpha$, the posterior uncertainty is reduced relative to the prior uncertainty. Along parameter directions that is less data-informed, i.e. $\sigma_i^2 < \alpha$, the posterior uncertainty increases relative to the prior uncertainty.

3. Applications to Imaging Problems.

3.1. Image Deblurring. One typical inverse problem in imaging is image deblurring. Given some blurry image, we want to recover the true, sharp image. To understand the deblurring process, we must first understand how an image becomes blurred in the first place. A simple and effective mathematical model of the blurring process is convolution of a sharp image with a blurring kernel. This blurring kernel is often described mathematically as a point spread function (PSF). The PSF describes how energy from a point source (ie. a single pixel) is *smear*ed out among neighboring pixels, resulting in a blur.

Since convolution is a linear operation, it can be expressed mathematically as

$$(3.1) \quad \mathcal{A}\mathbf{X}_{\text{true}} = \mathbf{B}$$

where \mathcal{A} is the blurring (convolution) operator acting on the true image $\mathbf{X}_{\text{true}} \in \mathbb{R}^{m_1 \times m_2}$ resulting in the blurred image $\mathbf{B} \in \mathbb{R}^{m_1 \times m_2}$. By stacking (or *vectorizing*) the columns of \mathbf{X}_{true} , we can write (3.1) as a linear algebraic equation. Let us denote by \mathbf{x}_{true} the vectorized true image and by \mathbf{y} the vectorized blurred image, i.e.,

$$\mathbf{x}_{\text{true}} = \text{vec}(\mathbf{X}_{\text{true}}) \in \mathbb{R}^{m_1 m_2}, \quad \mathbf{y} = \text{vec}(\mathbf{B}) \in \mathbb{R}^{m_1 m_2}$$

Also, since \mathcal{A} is a linear operator acting on a vector, it has a matrix representation denoted by $\mathbf{A} \in \mathbb{R}^{m_1 m_2 \times m_1 m_2}$. Finally, (3.1) becomes

$$(3.2) \quad \mathbf{A}\mathbf{x}_{\text{true}} = \mathbf{y}$$

Note that while this notation is convenient for manipulating mathematically, it is not efficient to construct the two-dimensional convolution matrix. \mathbf{A} is a large sparse matrix, which, for large problems, cannot be stored in memory. Even on problems small enough to fit in memory, it is computationally expensive to explicitly construct this matrix. Fortunately, there are efficient methods for computing spectral decompositions of the matrices arising from convolution operators using the fast fourier transform and discrete cosine transform. While interesting in their own right, these implementation details are not necessary for the following discussion. For a detailed treatment of image deblurring problems and algorithms, the interested reader is encouraged to consult [5].

For all examples considered in this chapter

$$\mathbf{\Lambda} = \lambda^2 \mathbf{I}, \quad \text{and } \mathbf{\Gamma} = \mathbf{I},$$

where λ is the noise level (the standard deviation).

Since truncated SVD (TSVD) and Tikhonov are spectral filtering methods, the regularized solution using these methods can be written using the following common form

$$(3.3) \quad \mathbf{x}_{filt} = \sum_{i=1}^p \phi_i \frac{\mathbf{U}_i^T \mathbf{y}}{\sigma_i} \mathbf{V}_i,$$

where ϕ_i is usually called the *filter factor* as it has the effect of filtering (damping) when ϕ_i is close to 0. It can be shown that the filter factor for rank- r TSVD is given by

$$\phi_i = \begin{cases} 1, & i \leq r \\ 0, & \text{otherwise.} \end{cases}$$

Likewise, the filter factor for Tikhonov regularization is given by

$$\phi_i = \frac{\sigma_i^2}{\sigma_i^2 + \alpha}$$

As discussed in Section 2, the DI method with rank- r approximation removes regularization on the first r directions \mathbf{V}_i , $1 \leq i \leq r$, while being the same as Tikhonov on the other directions. For $\mathbf{\Gamma} = \mathbf{I}$ and $\mathbf{x}_0 = 0$ the DI solution (see Lemma 2.2) can be written in the filtered form as

$$\phi_i = \begin{cases} 1, & i \leq r \\ \frac{\sigma_i^2}{\sigma_i^2 + \alpha}, & \text{otherwise.} \end{cases}$$

Remark 3.1. It should be emphasized that the DI method also shares the same spectral decomposition form in this case because $\mathbf{\Gamma} = \mathbf{I}$ and $\mathbf{x}_0 = 0$. When $\mathbf{\Gamma} \neq \mathbf{I}$, singular vectors of $\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A}$ do not necessarily diagonalize both \mathbf{A} and $\mathbf{\Gamma}$ simultaneously. In other words, the filtered form (3.3) is not valid for the DI approach unless \mathbf{U} and \mathbf{V} are singular vectors of $\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}}$ and $\mathbf{x}_0 = 0$. When $\mathbf{x}_0 \neq 0$, there is an additional term contributed from \mathbf{x}_0 as shown in the DI solution given in Lemma 2.2.

We can see here again that 1) when $r \rightarrow 0$, DI approaches Tikhonov; 2) when $\alpha \ll \sigma_i$ for $i \leq r$, Tikhonov is close to DI; and 3) when $\alpha \rightarrow \infty$, DI converges to TSVD. This can be clearly seen in Figure 5(a) for a deblurring problem in which we plot the relative error between the deblurred images and the original ones for $m_1 = m_2 = 128$, $\lambda = 0.01$, $r = 400$, and a wide range of α . For the under-regularization regime, i.e. $\alpha < 1$, which should be avoided, the regularization is not sufficient to suppress the oscillations due to the high frequency modes for both Tikhonov and DI methods, resulting in inaccurate reconstructions. For reasonable-to-over regularization regimes, i.e. $\alpha > 1$, DI is the best compared to both Tikhonov and TSVD method as it combines the advantages from both sides. That is: 1) DI behaves similar to Tikhonov for reasonable (but small) regularization and outperforms Tikhonov in reasonable-to-over regularization regimes; and 2) Compared to TSVD, DI is more accurate for reasonable regularization parameters as it maintains the benefits of keeping useful information from all parameter directions while avoiding potential errors caused by over-regularization. Consequently, the DI error is the smallest of the three methods discussed for all $\alpha > 10^3$ and DI is robust with respect to the regularization parameter.

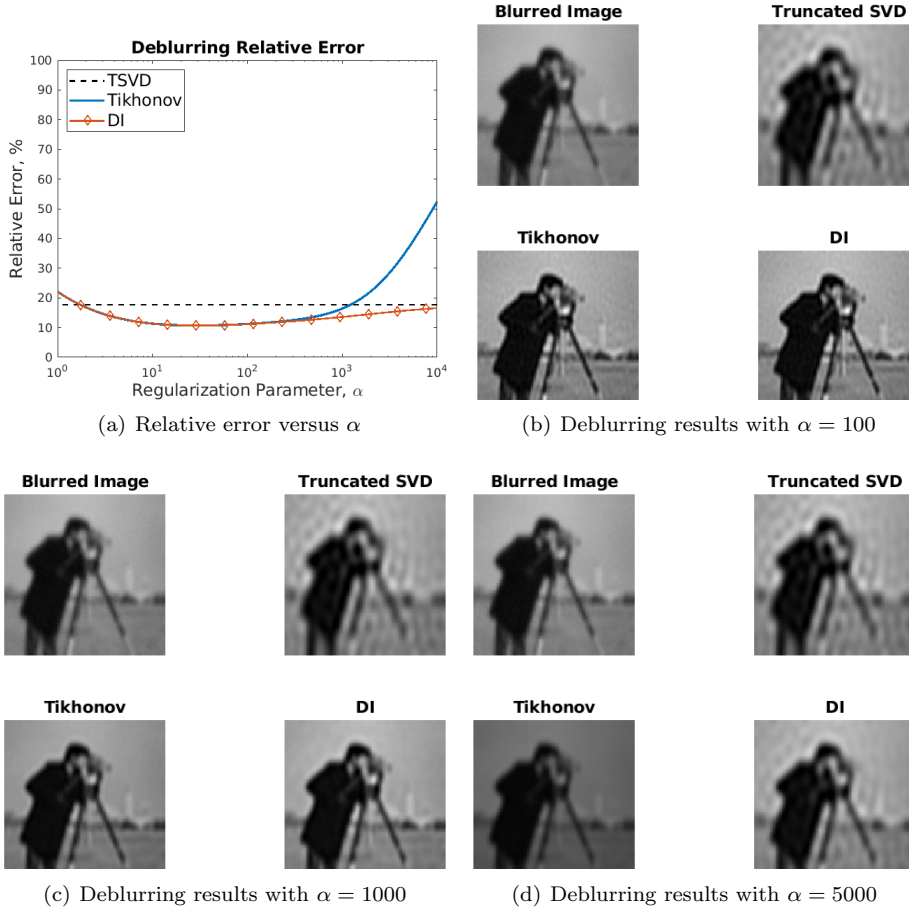


FIG. 5. Deblurring results for $m_1 = m_2 = 128$, $\lambda = 0.01$, $r = 400$. Top Left: relative error between deblurred images and the truth for a range of regularization parameter $\alpha \in [1, 10^4]$. Top Right: the DI deblurred image with $\alpha = 100$. Bottom Left: the DI deblurred image with $\alpha = 1000$. Bottom Right: the DI deblurred image with $\alpha = 5000$.

In Figure 5(b) are the deblurred images for $\alpha = 100$ corresponding to the smallest deblurring error for both DI and Tikhonov. As can be seen, Tikhonov result is similar to the DI one, while truncated SVD result is blurry as it removes (putting infinite regularization on) useful information in directions \mathbf{V}_i for $i > r$. Figures 5(c) and (d) show the deblurred images for $\alpha = 1000$ and $\alpha = 5000$, respectively, corresponding to cases where DI outperforms both Tikhonov and TSVD (see Figure 5(a)). Indeed, the DI deblurred image has higher quality.

In order to see if the DI method is sensitive to noise, we now consider the case with $\lambda = 5\%$ noise. Deblurring accuracy for this case (purple) is shown in Figure 6(a) together with the accuracy for the case of 1% noise (yellow). As can be seen, the solution quality of the DI method does not degrade significantly due to the presence of noise. Compare this to the difference seen in the Tikhonov method (red and blue curves) with the increase in noise level, we can see that the solution quality of the Tikhonov method degrades rapidly in the presence of noise. It can also be seen that Tikhonov regularization becomes more sensitive to the choice of α as the

noise increases. Since the DI method regularizes only the data-uninformed directions, which also contain much of the noise, increasing the noise level has little effect on the solution quality.

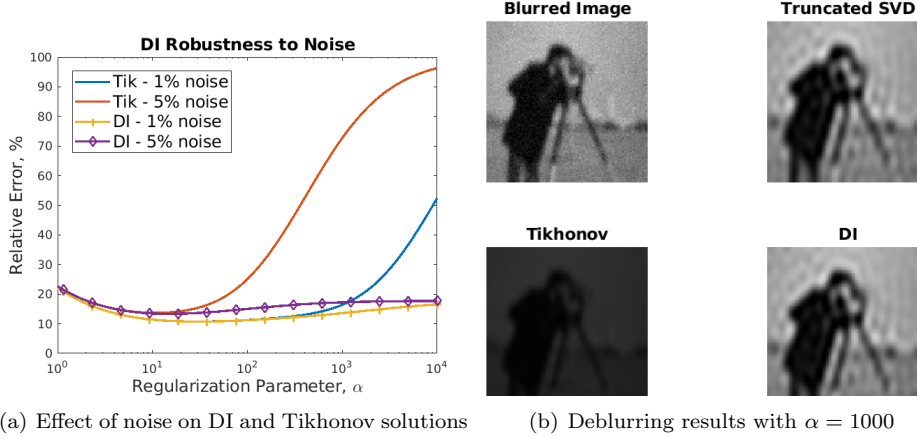


FIG. 6. Deblurring results for $m_1 = m_2 = 128$, $\lambda = 0.05$, $r = 400$. Left: relative error of DI and Tikhonov solutions with respect to true solution for noise levels of 1% and 5% and $\alpha \in [1, 10^4]$. Right: the DI, Tikhonov, and TSVD deblurred images with $\alpha = 1000$.

For the rest of this section, we consider the more challenging cases with $\lambda = 5\%$ noise. To make the problem even more challenging, we consider images with missing pixels to simulate more interesting cases when images are damaged or incomplete. Figure 7 show the deblurring results using DI, TSVD, and Tikhonov (Tik) regularizations for damaged images with $m_1 = m_2 = 128$, $r = 400$. First column contains four scenarios with 10% random data, 25% random data, 50% random data, and 100% data, all with noise. Note that we plot the damaged images by filling the missing data with 0. The second column are the corresponding TSVD deblurring results. The last four columns are the results from DI and Tikhonov with $\alpha = 10$ and 20. As can be observed, all methods are able to deblur and at the same time recover the true image quite well even with only 10% data. Both DI and Tikhonov yield clearer images compared to TSVD. The Tikhonov results are "darker", especially with $\alpha = 20$, indicating over-regularization, while the DI images are insensitive to regularization parameter as the data-informed modes are left untouched. Indeed, Figure 8 clearly demonstrates these expected results for larger regularization parameters ($\alpha = 50$ and $\alpha = 100$).

Recall the goal of sections 2.2 and 2.3.2 is to gain insights into statistical properties of the DI prior. For linear parameter-to-observable maps—which are the cases for this chapter—with Gaussian observational noise, the posterior is also a Gaussian. As a result, the result at the end of section 2.3.2 also allows us to use the posterior covariances (2.27) and (2.28) to estimate the uncertainty in the corresponding inverse solutions. Since the posterior using either Tikhonov or DI prior is Gaussian, its diagonal contains the marginal pixel-wise variances, which can be used as a measure of uncertainty for each pixel. We now study the uncertainty estimation in the solution of deblurring problems.

To begin, it is important to distinguish the following two cases:

- *Case I*: using only rank- r DI regularization in which rank- r approximation for




FIG. 7. Deblurring results using DI, TSVD, and Tikhonov (Tik) regularizations for damaged images with $m_1 = m_2 = 128$, $\lambda = 0.05$, $r = 400$. First column consists of four scenarios with 10%, 25%, 50%, and 100% data. The second column are the corresponding TSVD deblurring results. The last four columns are the results from DI and Tikhonov with $\alpha = 10$ and 20.

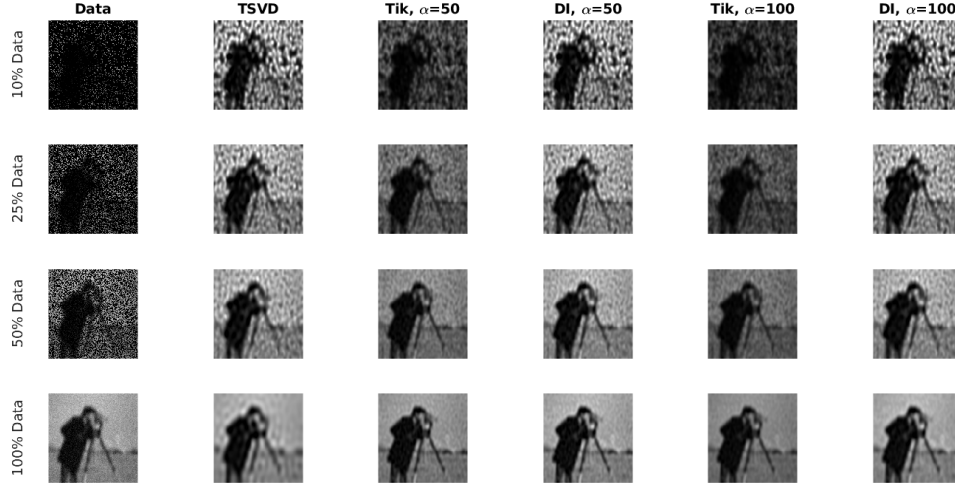


FIG. 8. Deblurring results using DI, TSVD, and Tikhonov (Tik) regularizations for damaged images with $m_1 = m_2 = 128$, $\lambda = 0.05$, $r = 400$. First column consists of four scenarios with 10%, 25%, 50%, and 100% data. The second column are the corresponding TSVD deblurring results. The last four columns are the results from DI and Tikhonov with $\alpha = 50$ and 100.

the pseudo-inverse $\left(\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma} \mathbf{A}^T \mathbf{\Lambda}^{-\frac{1}{2}}\right)^\dagger$ is done as we have presented. The DI posterior covariance (2.27) thus involves second and third columns in Table 1 and a rank- n SVD (2.13) is needed.

- *Case II:* performing rank- r low-rank approximation of the posterior covariance in addition to rank- r DI regularization. This amounts to using only the

second column of Table 1 for the DI posterior covariance in (2.27). This case is typically more practical for large-scale problems as only rank- r SVD (2.19) is needed.

In Figure 9(a) are the minimum pixel-wise variances for four scenarios with 10% random data, 25% random data, 50% random data, and 100% random data for *Case II*. As can be seen, the uncertainty corresponding to the case of missing data is lower than the uncertainty for full data case! We expect the opposite, that is, more available (supposedly) informative data is expected to lead to lower uncertainty in the inverse solution. The observation is twofold: first, care needs to be taken for *Case II* results as rank- r approximation may not provide accurate uncertainty; second, for 10% data case, when $r > 500$ the uncertainty is larger compared to the full data case. This suggests that r needs to be sufficiently large for an accurate uncertainty estimation, and this will be confirmed in the discussion below for *Case I* in which we use the full rank (rank- n) decomposition (2.13). The criteria for estimating such a value of r is a subject for future research.⁴

We next discuss the results for *Case I*. Again, this requires a rank- n SVD decomposition (2.13), where n is the rank of \mathbf{A} , to compute (2.27) using Table 1. Figure 9(b) shows that the minimum uncertainty for any missing data case is higher than the full data case regardless of any value of r in rank- r DI regularization. As also expected, the uncertainty scales inversely with the amount of available data, i.e., the more informative data we have the smaller the uncertainty in the inverse solution. Note that the result and the conclusion for the largest pixel-wise variances are similar and hence omitted here.




FIG. 9. Rank- r DI posterior pixel-wise uncertainty using rank- n SVD decomposition (*Case I* with both second and third columns of Table 1) and using rank- r SVD decomposition (*Case II* with only the second column of Table 1).

We now compare the DI and Tikhonov posterior uncertainty estimations. Since *Case I*, though more expensive, provides more accurate uncertainty estimation, it is used for computing DI posterior pixel-wise variances. To be fair, we also use the full decomposition for Tikhonov regularization. In other words, the following comparison is based on (2.27) and (2.28) and Table 1. As discussed above in Figure 6(a) and Figure 7, $\alpha = 10$ corresponds to a case in the region where DI and Tikhonov give

⁴At the moment of writing this chapter we have not yet found such a criteria.

nearly the same reconstructions (in fact Tikhonov slightly over-regularizes), so let us start with this case first. Figure 10 shows that the DI posterior has higher pixel-wise variance than the Tikhonov posterior, this is consistent with the result and the discussion of Table 1 and Figure 7, that is, the Tikhonov posterior is not only over-regularizing but also overconfident. For both methods, regions of higher uncertainty are visually discernible where data is missing. In the case of 100% data, the result is the same, namely, Tikhonov uncertainty estimation subjectively is less than the DI uncertainty estimation. In this case, the uncertainty estimate is not very interesting: both DI and Tikhonov have approximately uniform uncertainty everywhere as we have data everywhere. We next consider the case with $\alpha = 1000$ where Tikhonov significantly over-regularizes (see Figure 6(b)). Figure 11, shows that while Tikhonov is uniformly (very) overconfident, i.e., having small posterior uncertainty everywhere, DI gives informative UQ results. The latter can be clearly seen for the case with 10% data in which the uncertainty is higher for missing pixels. This implies that the DI priors could provide more useful UQ results than the Tikhonov (standard Gaussian) ones.

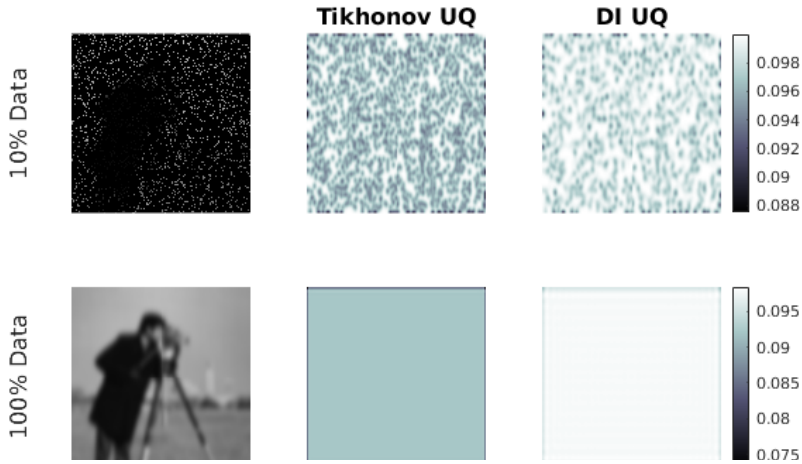


FIG. 10. Visualization of pixel-wise variance estimates for the deblurring problem with $\lambda = 0.05$, $r = 400$, and $\alpha = 10$. In the left column are the noisy images with 10% data and 100% data. In the second column are the Tikhonov uncertainty estimates for 10% data (top) and 100% data (bottom). Likewise, the third column contains the DI uncertainty estimates for 10% data (top) and 100% data (bottom).

3.2. Image Denoising. We can extend the idea of data-informed (DI) regularization to the image denoising problem. Since noise typically resides in the high frequency portion of the image, denoising can be performed by applying spectral filtering techniques directly to the noisy image. These noisy high-frequency modes are also the *less informative* modes in the DI setting. Taking the SVD of the noisy image, \mathbf{X}_{noisy} , we have

$$\mathbf{X}_{noisy} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_i \sigma_i \mathbf{U}_i \mathbf{V}_i^T,$$




FIG. 11. Visualization of pixel-wise variances for the deblurring problem with $\lambda = 0.05$, $r = 400$, and $\alpha = 1000$. In the left column are the noisy images with 10% data and 100% data. In the second column are the Tikhonov uncertainty estimates for 10% data (top) and 100% data (bottom). The third column contains the DI uncertainty estimates for 10% data (top) and 100% data (bottom).

The denoised image can be obtained by "filtering" the noise as

$$\mathbf{X}_{filt} = \mathbf{U} \mathbf{\Sigma}^{filt} \mathbf{V}^T = \sum_i \phi_i \sigma_i \mathbf{U}_i \mathbf{V}_i^T,$$

where $\mathbf{\Sigma}^{filt}$ is the diagonal matrix with $\Sigma_{ii}^{filt} = \phi_i \sigma_i$. The filter factors ϕ_i are the same as those defined for the deblurring case. For a numerical demonstration, we pick a noisy image [5] with 5% noise (see the top left sub-figure of Figure 12(a)). Shown in Figure 12(a) are denoised results using DI with $r = 20$ and $\alpha = 100$, TSVD with $r = 20$, and Tikhonov with $\alpha = 100$. Though the difference in the results is not clearly visible, the DI has smaller error compared the other two methods. This can be verified in Figure 12(b) where the relative error between the denoised image and the true one for a wide range of "regularization parameter" $\alpha \in [10^{-2}, 10^4]$ is presented. Clearly, we would not choose $\alpha < 1$ as these correspond to under-regularization. For $\alpha > 1$, DI is the best compared to both Tikhonov and TSVD method as it combines the advantages from both methods. Indeed, the DI error is smallest for all $\alpha > 1$ and DI is robust with regularization parameter.

3.3. X-ray Tomography. In the previous two examples, we have been able to implement spectral filtering methods directly by introducing filter factors which effectively modified the singular values to minimize the impact of noise on the inversion process.⁵ Each method relied on computing a full factorization of $\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A}$ and then applying filters. While this is an effective and straightforward method to solve small-to-moderate inverse problems that helps provide insight into each approach, it can be cumbersome or even computationally infeasible to compute full factorizations for large-scale problems. It is not uncommon that inverse problems arising in imaging applications can lead to very large matrix operators. Indeed, we have seen even in

⁵Recall that the DI method also shares the same spectral decomposition form in this case because $\mathbf{\Gamma} = \mathbf{I}$ and $\mathbf{x}_0 = 0$.

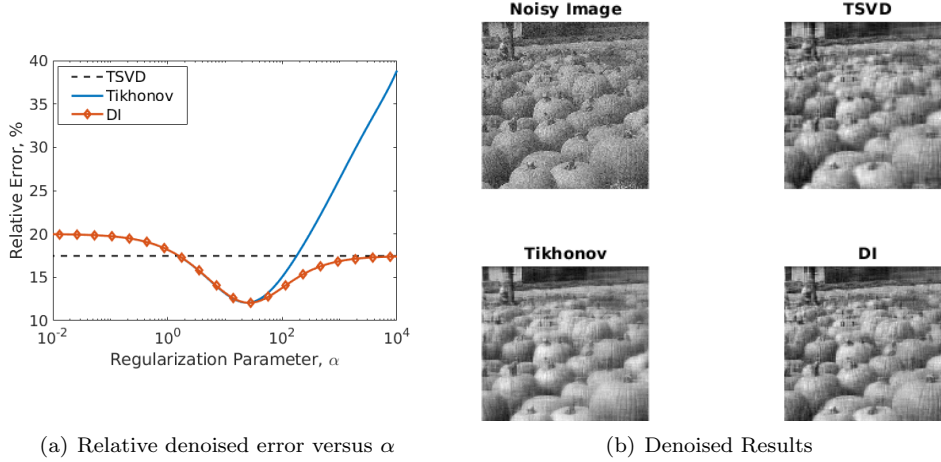


FIG. 12. Denoising with DI, Tikhonov, and TSVD methods. Left: The relative error between the denoised image and the true one for a wide range of "regularization parameter" $\alpha \in [10^{-2}, 10^4]$. The DI error is smallest for all $\alpha > 1$ (corresponding reasonable to over-regularization regimes). Right: Denoised results using DI with $r = 20$ and $\alpha = 100$, TSVD with $r = 20$, and Tikhonov with $\alpha = 100$.

the toy image deblurring problem in Section 3.1 that matrix size of 16384×16384 is significantly large and we have employed more sophisticated methods to compute the factorization of the convolution operator. For many problems, however, such efficient factorizations may not exist or it is computationally prohibitive to compute a full factorization.

One way to overcome the challenge of factorizing large matrices is to solve the optimality condition (2.18) iteratively. Since \mathbf{H} is symmetric positive definite, we choose the conjugate gradient (CG) method (see, e.g., [11] and the references therein) which requires only matrix-vector products, which in turn avoids forming any matrices (including \mathbf{A} or \mathbf{H}) completely. We consider two variants: a) Using CG to solve for (2.20), that is, we still require rank- r approximation of the DI regularization; and b) Using CG to solve for (2.18), that is, a rank- r approximation of the DI regularization is not required and in this case we use a least squares optimization method to compute the pseudo-inverse $\left(\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma} \mathbf{A}^T \mathbf{\Lambda}^{-\frac{1}{2}}\right)^\dagger$ acting on a vector for each CG iteration.

The detailed computational procedure for the a)-variant is given in Algorithm 3.1 To demonstrate the effectiveness of this approach for the DI method, we choose to solve the inverse problem of reconstructing an image from X-ray measurements. The forward model of generating X-ray measurements, \mathbf{A} , is given by the Radon transform and \mathbf{A}^T is given by the inverse Radon transform. A more detailed description of the X-ray tomography inverse problem is given in [9] (and the references therein). We use the MATLAB Image Processing Toolbox to compute the product of the Radon transform \mathbf{A} and its inverse \mathbf{A}^T with a vector. Results using Algorithm 3.1 for a popular 256×256 phantom image are shown in Figure 13 for various values of the regularization parameter α and the rank r . Each row contains the results for each regularization parameter with different values of r . The corresponding values for α and r can also be found in the rows and columns of Table 2. Note that below each figure is the relative error of the corresponding reconstruction and the actual phantom

Algorithm 3.1 Data-informed Inversion Using Randomized Eigensolver and CG

Input: Data \mathbf{y} , number of eigenvectors r , prior \mathbf{x}_0 , prior covariance matrix $\mathbf{\Gamma}$, noise covariance matrix $\mathbf{\Lambda}$, regularization parameter α

- 1: Define $\mathbf{F} := \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}}$.
- 2: Create functions to compute matrix-vector products $\mathbf{F}\mathbf{x}$ and $\mathbf{F}^T\mathbf{x}$
- 3: Compute the first r eigenvectors of $\mathbf{F}^T\mathbf{F}$ using a randomized eigensolver
- 4: Solve linear equation (2.20), i.e.,

$$\mathbf{\Gamma}^{-\frac{1}{2}} \left[\mathbf{F}^T \mathbf{F} + \alpha (\mathbf{I} - \mathbf{V}_r \mathbf{V}_r^T) \right] \mathbf{\Gamma}^{-\frac{1}{2}} \mathbf{x} = \mathbf{\Gamma}^{-\frac{1}{2}} \mathbf{F}^T \mathbf{y} + \alpha \mathbf{\Gamma}^{-\frac{1}{2}} (\mathbf{I} - \mathbf{V}_r \mathbf{V}_r^T) \mathbf{\Gamma}^{-\frac{1}{2}} \mathbf{x}_0$$

using the conjugate gradient method

image. These relative errors are collected in Table 2 for clarity. Note that for the last two images on the last row of Figure 13, CG does not converge and this issue is still under investigation. Other than that the observations are similar to the previous section. That is, compared to Tikhonov, DI is robust to the regularization parameter and it is at least as good as Tikhonov regardless the values of regularization parameter α and rank r .

α	Relative Error, %					
	$r = 0$ (Tik)	$r = 10$	$r = 50$	$r = 100$	$r = 200$	$r = 400$
1	33.52	33.52	33.52	33.52	33.52	33.52
10	31.73	31.73	31.73	31.73	31.73	31.73
100	24.44	24.45	24.45	24.45	24.45	24.45
1000	29.81	29.80	29.72	29.66	29.51	29.09
10^4	58.76	58.52	56.93	55.92	54.03	50.43
10^5	81.77	77.10	70.33	67.78	63.84	57.84
10^6	96.09	81.29	72.44	69.50	81.80	299.73

TABLE 2

Comparison of the relative errors of the DI solution estimate for various regularization parameters α and various values for r . The noise level here is $\lambda = 1\%$.

Next we present the detailed computational procedure for the b)-variant in Algorithm 3.2. In order to compare variant b) with variant a), we compute the relative error of the reconstruction and the true image for various values of regularization parameter α . From the results in Figure 13, we choose $r = 200$ to balance the accuracy and the cost of the eigensolver. The result is in Figure 14, which shows that the b)-variant (red curve) is at least as good as the a)-variant (blue curve) while not requiring low rank approximations. Indeed, to demonstrate this, we pick $\alpha = 100$ for which Figure 14 shows that both variants give similar reconstruction quality, and the reconstruction from both variants are shown in Figure 15. As can be seen, the result from the b)-variant looks much clearer, which is expected in this case, as $r = 200$ is not sufficient to capture all the data-informed modes for the a)-variant.

By using the pseudoinverse formulation, we can still get excellent results while avoiding the computation of a large factorization.

4. Conclusions. We have presented a new regularization technique called data-informed (DI) regularization that combines advantages of the classical truncated SVD




FIG. 13. *DI reconstructions for various values of the regularization parameter α and the rank r . Each row contains the results for each regularization parameter with different values of r . The corresponding values for α and r can be found in the rows and columns of Table 2. Below each figure is the relative error of the corresponding reconstruction and the actual phantom image.*

and Tikhonov regularization. In particular, the DI approach does not pollute the data-informed modes, and regularizes only less data-informed ones. As a direct consequence, the DI approach is at least as good as the Tikhonov method for any value of the regularization parameter and it is more accurate than the TSVD (for reasonable regularization parameter). Due to the blending of these two classical methods, DI is expected to be robust with regularization parameter and this is verified numerically. We have shown that our DI approach has an interesting statistical interpretation, that is, it transforms both the data distribution (i.e. the likelihood) and prior distribution (induced by Tikhonov regularization) to the same Gaussian distribution whose covariance matrix is diagonal and the diagonal elements are exactly the singular values of a composition of the prior covariance matrix, the forward map, and the noise covari-

Algorithm 3.2 Data-informed Inversion Using Nested CG

Input: Data \mathbf{y} , number of eigenvectors r , prior \mathbf{x}_0 , prior covariance matrix $\mathbf{\Gamma}$, noise covariance matrix $\mathbf{\Lambda}$, regularization parameter α

- 1: Define $\mathbf{F} = \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}}$
- 2: Create functions to compute matrix-vector products $\mathbf{F}\mathbf{x}$ and $\mathbf{F}^T\mathbf{x}$
- 3: Solve linear equation (2.18), i.e.,

$$\mathbf{\Gamma}^{-\frac{1}{2}} \left[\mathbf{F}^T \mathbf{F} + \alpha (\mathbf{I} - \mathbf{F}^T (\mathbf{F} \mathbf{F}^T)^\dagger \mathbf{F}) \right] \mathbf{\Gamma}^{-\frac{1}{2}} \mathbf{x} = \mathbf{\Gamma}^{-\frac{1}{2}} \mathbf{F}^T \mathbf{y} + \alpha \mathbf{\Gamma}^{-\frac{1}{2}} (\mathbf{I} - \mathbf{F}^T (\mathbf{F} \mathbf{F}^T)^\dagger \mathbf{F}) \mathbf{\Gamma}^{-\frac{1}{2}} \mathbf{x}_0$$

using the conjugate gradient method. For each CG iterations, compute the product of $\mathbf{F}^T (\mathbf{F} \mathbf{F}^T)^\dagger \mathbf{F} \mathbf{\Gamma}^{-\frac{1}{2}}$ with any vector \mathbf{x} using matrix-free Algorithm 3.3.

Algorithm 3.3 Compute the product of $\mathbf{F}^T (\mathbf{F} \mathbf{F}^T)^\dagger \mathbf{F} \mathbf{\Gamma}^{-\frac{1}{2}}$ with any vector using optimization

Input: functions to compute $\mathbf{F}\mathbf{x}$ and $\mathbf{F}^T\mathbf{x}$, current estimate of \mathbf{x} , prior covariance matrix $\mathbf{\Gamma}$

- 1: Compute $\mathbf{b} = \mathbf{F} \mathbf{\Gamma}^{-\frac{1}{2}} \mathbf{x}$
- 2: Using conjugate gradient method, solve linear equation

$$\mathbf{F} \mathbf{F}^T \mathbf{z} = \mathbf{b}.$$

- 3: Return $\mathbf{F}^T \mathbf{z}$





FIG. 14. A comparison between variant b) (red curve) and variant a) with $r = 200$ (blue curve). Here, we compute the relative error of the reconstruction and the truth image for various values of regularization parameter α .

ance matrix. In other words, DI finds the modes that are most equally data-informed and prior-informed and leaves these modes untouched so that the inverse solution receive the best possible (balanced) information from both prior and the data. We



(a) X-ray tomography using a)-variant with $r = 200$ (b) X-ray tomography using b)-variant

FIG. 15. X-Ray tomography reconstruction with 1% noise and $\alpha = 100$: (a) the result from the a)-variant with $r = 200$, and (b) the result from the b)-variant.

have shown that DI is a regularization strategy. To demonstrate and to support our findings, we have presented various results for popular computer vision and imaging problems including deblurring, denosing, and X-ray tomography.

Acknowledgments. This work is inspired by an ongoing collaborative work on data-consistent inverse framework with Tim Wildey and Brad Marvin, and the authors would like to thank them for many fruitful discussions.

REFERENCES

- [1] A. C. ANTOUNAS, *Approximation of Large-Scale Systems*, SIAM, Philadelphia, 2005.
- [2] D. COLTON AND R. KRESS, *Integral equation methods in scattering theory*, John Wiley & Sons, 1983.
- [3] J. N. FRANKLIN, *Well-posed stochastic extensions of ill-posed linear problems*, Journal of Mathematical Analysis and Applications, 31 (1970), pp. 682–716.
- [4] S. GUGERCIN AND A. C. ANTOUNAS, *A survey of model reduction by balanced truncation and some new results*, International Journal of Control, 77 (2004), pp. 748–766.
- [5] N. J. G. HANSEN, P. C. AND D. P. O’LEARY, *Deblurring Images: Matrices, Spectra, and Filtering*, SIAM, Philadelphia, 2006.
- [6] J. KAIPIO AND E. SOMERSALO, *Statistical and Computational Inverse Problems*, vol. 160 of Applied Mathematical Sciences, Springer-Verlag, New York, 2005.
- [7] S. LASANEN, *Discretizations of generalized random variables with applications to inverse problems*, PhD thesis, University of Oulu, 2002.
- [8] M. S. LEHTINEN, L. PÄIVÄRINTA, AND E. SOMERSALO, *Linear inverse problems for generalized random variables*, Inverse Problems, 5 (1989), pp. 599–612.
- [9] S. S. MUELLER, J. L., *Linear and Nonlinear Inverse Problems with Practical Applications*, SIAM, Philadelphia, 2012.
- [10] P. PIHOINEN, *Statistical measurements, experiments, and applications*, PhD thesis, Department of Mathematics and Statistics, University of Helsinki, 2005.
- [11] J. R. SHEWCHUK, *An introduction of the conjugate gradient method without the agonizing pain*, (1994), <https://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.pdf>.
- [12] A. M. STUART, *Inverse problems: A Bayesian perspective*, Acta Numerica, 19 (2010), pp. 451–559, <https://doi.org/doi:10.1017/S0962492910000061>.
- [13] A. TARANTOLA, *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM, Philadelphia, PA, 2005.