

RED-PSM: Regularization by Denoising of Factorized Low Rank Models for Dynamic Imaging

Berk Iskender , Graduate Student Member, IEEE, Marc L. Klasky , and Yoram Bresler , Life Fellow, IEEE

Abstract—Dynamic imaging addresses the recovery of a time-varying 2D or 3D object at each time instant using its under-sampled measurements. In particular, in the case of dynamic tomography, only a single projection at a single view angle may be available at a time, making the problem severely ill-posed. We propose an approach, RED-PSM, which combines for the first time two powerful techniques to address this challenging imaging problem. The first, are non-parametric factorized low rank models, also known as partially separable models (PSMs), which have been used to efficiently introduce a low-rank prior for the spatio-temporal object. The second is the recent *Regularization by Denoising (RED)*, which provides a flexible framework to exploit the impressive performance of state-of-the-art image denoising algorithms, for various inverse problems. We propose a partially separable objective with RED and a computationally efficient and scalable optimization scheme with variable splitting and ADMM. Theoretical analysis proves the convergence of our objective to a value corresponding to a stationary point satisfying the first-order optimality conditions. Convergence is accelerated by a particular projection-domain-based initialization. We demonstrate the performance and computational improvements of our proposed RED-PSM with a learned image denoiser by comparing it to a recent deep-prior-based method known as TD-DIP. Although the main focus is on dynamic tomography, we also show performance advantages of RED-PSM in a cardiac dynamic MRI setting.

Index Terms—Dynamic imaging, regularization by denoising, low rank modeling.

I. INTRODUCTION

TIME-VARYING or dynamic tomography involves the reconstruction of a dynamic object using its projections acquired sequentially in time. The problem arises in micro tomography [1], myocardial perfusion imaging [2], thoracic CT [3], imaging of fluid flow processes [4], [5] and dynamic imaging of material samples undergoing compression [6], [7]. Also, it is closely related to the dynamic MRI (dMRI) problem, which typically arises in cardiac imaging [8].

Dynamic tomography is a challenging ill-posed inverse problem: since the measurements are inconsistent due to the evolving

object, traditional reconstruction algorithms lead to significant artifacts. As discussed further below, although numerous methods have been proposed to address this problem, they suffer from various limitations, or provide less than satisfactory reconstruction quality, especially in the scenario of main interest in this paper, when at each time instant only one projection is available.

A. Proposed Approach

We propose an approach, RED-PSM, which combines for the first time two powerful techniques to address this challenging imaging problem. The first technique, are non-parametric factorized low-rank object models, popularized under the name *partially separable models (PSMs)* [9]. PSM was combined with low-rank matrix recovery in [10], [11], [12], and used since then to represent or motivate a *low-rank object prior* in tens of works on dynamic imaging. The second technique is the recent *Regularization by Denoising (RED)* [13], which provides a flexible framework to exploit the impressive performance of state-of-the-art image denoising algorithms, for various inverse problems.

We propose a non-convex partially separable objective with RED for learning-regularized low-rank spatio-temporal object recovery and a computationally efficient and scalable optimization scheme with bi-convex ADMM.

Theoretical analysis proves the convergence of our objective to a value corresponding to a stationary point satisfying the first-order optimality conditions. Convergence is accelerated by a particular projection-domain-based initialization.

We demonstrate the performance and computational improvements of our proposed RED-PSM with a learned image denoiser by comparing it to a recent deep-prior-based method known as TD-DIP [14], and to spatial and spatiotemporal total variation (TV) regularized versions of PSM with a bi-convex objective.

Although the main focus is on dynamic tomography, we also show the performance advantages of RED-PSM in a cardiac dynamic MRI setting.

RED was originally inspired by the Plug-and-Play (PnP) approach [15], which was the first to exploit powerful denoisers to represent the prior on the object in inverse problems but without explicitly defining the regularizer. Instead, in PnP the denoiser is incorporated by replacing a proximal mapping. Both PnP and RED have shown strong empirical performance in various applications. A recent survey [16] provides a detailed discussion of the main results and different applications of PnP and the relation to RED. For the algorithmic variations PnP-ADMM and PnP-ISTA there are convergence results to unique fixed points. Moreover, for PnP-ISTA, similar to our result, convergence is shown to a stationary point of a possibly non-convex objective

Manuscript received 2 August 2023; revised 20 November 2023, 13 February 2024, and 2 April 2024; accepted 5 May 2024. Date of publication 20 May 2024; date of current version 31 May 2024. This work was supported by Los Alamos National Labs under Subcontract 599416/CW13995. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Bo Zhao. (*Corresponding author: Berk Iskender.*)

Berk Iskender and Yoram Bresler are with the Department of Electrical and Computer Engineering, and the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: berki2@illinois.edu).

Marc L. Klasky is with the Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87544 USA.

This article has supplementary downloadable material available at <https://doi.org/10.1109/TCI.2024.3402347>, provided by the authors.

Digital Object Identifier 10.1109/TCI.2024.3402347

when the denoiser is an MMSE estimator [17]. This result requires the prior to be non-degenerate, i.e., not lie on a lower dimensional manifold, which, unfortunately, would be violated by the low-rank bilinear representation in our PSM. We believe that a PnP-based method similar to RED-PSM can be formulated for our problem, but a proof of convergence to a stationary point may not directly follow, and a different approach and analysis may be needed.

Our motivation for using the original RED formulation in RED-PSM rather than more recent formulations and analyses (e.g., as in [18] and the references therein) are its simple gradient expression, and explicit regularizer expression facilitating implementation and theoretical analysis. The objective is to demonstrate the large improvement in performance that this simple learned regularizer can bring to the PSM scheme, and the improvement over much more complex and slower state-of-the-art methods. Recent variations on RED such as in [18], and PnP methods (as covered in [16]) have improved over the original methods in reconstruction quality in various static imaging problems and/or in theoretical guarantees. However, theoretical guarantees for the case of a non-convex data fidelity term have been provided recently for only one such non-convex scenario - phase-retrieval [19]. As this is the first work to extend RED to spatio-temporal imaging with a PSM, and to provide a convergence guarantee for the method in the non-convex bilinear scenario, we chose to limit our attention in this paper to the original RED version [13], [20].

B. Previous Work

We limit our discussion to dynamic CT and dMRI, although similar or analogous methods have been or can be applied in other dynamic imaging modalities. The same can be said of the method proposed in this paper.

1) *Dynamic MRI (dMRI)*: In spite of major advances in hardware and signal acquisition methods, MRI has remained a relatively slow modality, posing significant challenges for dynamic imaging. This, coupled with the high flexibility in the acquisition strategy, has motivated extensive work on dMRI, with hundreds of papers on the subject. We highlight some key developments in the algorithm-based dMRI techniques.

Fundamentally, algorithmic approaches to dMRI are based on identifying and exploiting redundancy in the dynamic object, which enable its recovery from an incomplete, and usually sparse set of samples acquired in the joint k -space and time domain – the k - t -space. Many of the proposed methods have extensions to parallel imaging, which introduces additional such redundancy by additional hardware.

A large set of approaches rely on modeling the support of the object in the domain dual to the k - t -domain – the $(x$ - $f)$ domain (the space – temporal frequency domain). These approaches may be classified into three groups. The first involves generic modeling of the $(x$ - $f)$ support as in reduced field-of-view (FOV) techniques [21], [22], [23]. The second involves sparse and unknown support modeling in a transform domain, and sampling and recovery using the methods of compressed sensing [24], [25], [26], [27]. The third involves adaptive sparse support modeling [28], [29], [30], [31], where the support information is adapted to the dynamics of the imaged slice, and the theory of time-sequential sampling [32], [33] is used to design an optimal k - t sampling and reconstruction scheme. Other related approaches model prior information such as motion periodicity

and related harmonic structure [34], [35], or the object’s spatial and temporal correlation functions [36].

Since its introduction, the partially-separable model (PSM) [9], with a non-parametric factorized low-rank form separating the spatial and temporal structure, has been used extensively in many works on dMRI to represent the underlying spatio-temporal object. A group of methods recover the PSM in two steps: (a) estimate the temporal subspace using data (e.g., a navigator signal) sampled at low spatial, but high temporal resolution; and (b) estimate the spatial subspace using the entire acquired data set. Such methods include e.g., [9], [37], [38], [39], [40], with followup versions promoting sparsity of the object in a transform domain (e.g., [41]). Other PSM approaches use a low-rank matrix recovery formulation to estimate spatial and temporal components jointly (free of conditions or additional data acquisition enabling a separate reconstruction of the temporal subspace), e.g., [10], [11]), also with versions promoting sparsity [12], [42].

Manifold models have also been considered for dMRI. Recent such work [43] surveys previous related works, and proposes a new method that is related to PSM. The method partitions the dynamic object into temporal subsets each lying on a low-dimensional manifold, approximated by its tangent subspace. As a result, the object is represented by a temporally partitioned PSM. A simple sparsity penalty is used for the spatial representation. The spatial basis functions and linear transformations aligning the temporal subsets are determined by optimizing a non-convex objective using an ADMM algorithm. The subset temporal partitioning and temporal bases are estimated using an initial reconstruction obtained from training data. In some applications, such data may not be available.

Another set of methods, including [44], [45], [46], [47], promotes low-rankness of the entire spatio-temporal object matrix or of patches thereof [48], [49] implicitly by using the nuclear norm or Schatten- p functional with $p < 1$, often promoting sparsity in a transform domain as well. Methods combining motion estimation and compensation with other priors are also used widely for dMRI, e.g., [50], [51], [52], [53].

A different approach in dMRI decomposes the object into the sum of low-rank and sparse components, with the low-rankness encouraged implicitly using nuclear or a Schatten- p functional [54], [55], [56], [57]. A recent related fast method [58] decomposes the object representation into the sum of three components: the mean signal; a low-rank PSM; and a residual sparse in the Fourier domain. Although benchmark competing methods are either faster or more accurate on some of the diverse data sets used for the comparison in [58], the method provides the lowest reconstruction error and runtimes when these quantities are averaged over these data sets.

Recent approaches to dMRI reconstruction have used deep neural networks (DNNs). A majority of these methods, such as [59], [60], [61], [62], [63], [64], [65], [66], [67], [68], [69], [70], [71], [72], are supervised, requiring a large set of high-resolution artifact-free images for network training. However, acquiring such training data can be difficult for some dMRI applications, and may be impossible for many other dynamic imaging applications, where clinical dMRI strategies for freezing motion such as ECG synchronization and breath-hold by a cooperative patient, are not available.

The need for a training data set is overcome by recent object-domain deep image prior (DIP)-based DNN algorithms for

dMRI [14], [73], which are unsupervised. Providing impressive results, DIP-based algorithms such as [14] suffer from overfitting and usually require handcrafted early stopping criteria during the optimization of generator network parameters. To overcome these drawbacks, [73] includes regularization constraining the geodesic distances between generated objects. However, this requires the computation of the Jacobian of the generator network at each iteration of the update of the weights, significantly increasing the computation and run time.

A different unsupervised approach [74] combines factorized low rank (i.e., partially separable) with generative models. Although it combines the PSM with the recent DIP framework, this method has the following limitations: i) The spatial generator is an artifact removal network taking the full-sized spatial basis functions as input. Since this prevents a patch-based implementation, it may be computationally problematic for high-resolution 3D+temporal settings; ii) The CNN structural prior for natural images that is used in the spatial generator may not be useful as a prior for the individual spatial basis functions, since the least-squares optimal spatial basis functions are the left singular vectors of the complete object, and as such may not have the structure of natural images; iii) As with the other DIP-based methods, if the additional penalties on the generator parameters are insufficient, this method can be prone to overfitting. These various limitations are overcome by the proposed approach.

2) *Dynamic Tomography (dCT)*: CT is faster than MRI, thanks to acquiring an entire projection at one view angle at once. However, because of the angular scanning required in almost all systems, in spite of significant advances in hardware, dCT is still a relatively slow modality. It is also the problem that motivated much of our work reported in this paper.

In perhaps the earliest algorithmic approaches to dCT, the time-sequential sampling problem in d-CT was formulated and studied in [75], and [76], [77] provided a solution using time-sequential sampling of bandlimited spatio-temporal signals including an optimal view angle order for the scan and theoretical guarantees for unique and stable reconstruction. However, the approach is limited by its bandlimitedness assumptions.

A class of algorithms for dCT are based on modeling the time variation by a motion field, and motion-compensated reconstruction. Many of these algorithms, e.g., [78], [79], [80], [81] alternate between estimating the motion field and the time-varying object. Some algorithms, e.g., [82] use the optical flow (mass continuity) PDE. A different method [83] estimates the motion field separately by evolving a linear elastic deformation by the Navier-Cauchy PDE, assuming known boundary evolution of the object and initial density distribution. However, these methods assume the total mass or density to be preserved over time, which may be limiting assumptions. For instance, flow of a contrast agent into the imaged volume, or imaging a fixed slice of a 3D time-varying object under compression with cross-slice motion may violate either of these assumptions. As a hardware-based, problem-specific approach, optical tracking of fiducial marks has been used for motion correction [84].

Various methods approach dCT reconstruction as Bayesian estimation. These include methods [85], [86], [87], [88] that employ a state-space formulation and use Kalman filter techniques to approximate an MMSE estimate. A different Bayesian method [89] incorporates a spatiotemporal Markov random field object prior, and models measurement imperfections. An iterative algorithm is used to compute a MAP estimate. An interlaced projection acquisition scheme is also proposed.

Several dCT methods motivated by compressed sensing use a total-variation penalty along both spatial and temporal coordinates as in [90], while other methods, e.g., [91] propose efficient sparse representations for the dynamic object. Incorporating low-rank modeling [92] decomposes the object into low-rank+sparse components, promoting the low rankness implicitly using the nuclear norm. The method relies on the object being approximately static for groups of projections, enabling an approximate reconstruction from each such group.

The PSM model has been introduced into dynamic tomography [93], [94] by carrying the PSM to the projection domain using the harmonic representation of projections, and estimating the spatial and temporal basis functions jointly. The work provides a theoretical analysis of uniqueness and stability and the choice of time-sequential angular scan scheme for the problem. Despite the advantages of this approach, its performance is still limited by the null space of the measurement operator.

To help address the underdetermined problem, the method of [95] improves over [93], [94] by combining a PSM with basic spatial regularization by total variation. Like [93], [94] this method too estimates the temporal and spatial PSM components jointly. This makes it applicable to the time-sequential acquisition scenario considered in this paper, where only a single projection is available per time frame. It imposes low rankness as a soft constraint, using a hybrid biconvex objective with a data fidelity term and a TV regularizer expressed in terms of the unfactorized object, and a penalty for object-PSM mismatch. It uses [93], [94] for fast initialization. The method [95] can be considered a precursor to the method of this paper, which combines a hard rank constrained PSM with a sophisticated learned spatial regularizer, and provides theoretical convergence analysis.

Recent methods [96], [97], [98], [99] introduce deep learning into dCT. These methods require an initial reconstruction from the subsets of data from which the motion can be estimated accurately. However, without a periodicity assumption, this is not applicable to the scenarios in this paper, where only a single projection is available at a time. Another method [100] for cone-beam dCT proposes a multi-agent consensus equilibrium [101]-based technique using separate deep denoisers for axial, sagittal, and coronal slices of the object.

A deep-learning-based method [102] for dynamic photoacoustic tomography (PACT) introduced a neural field (NF) to represent the dynamic object with fewer parameters, combining it with a simple total variation (TV)-based regularizer. Despite providing an efficient representation for the dynamic object, the method outperforms its comparison benchmarks only in some of the projection-per-frame settings.

C. Contributions

1) To the best of our knowledge, RED-PSM is the first PSM-based approach to dynamic imaging that incorporates a particular (RED-based [13]) spatial prior that is learned and pre-trained.

2) We are not aware of any prior work that uses RED with an explicit low-rankness constraint. In RED-PSM we achieve parsimonious representation for the dynamic object using PSM, and reduce the data requirements further by incorporating the RED prior.

3) Unlike supervised learning-based methods for spatio-temporal imaging [59], [60], [61], [62], [63], [64], [65], [66],




Fig. 1. Imaging geometry for time-varying tomography of the object f_t with single measurement g_t at each time instant for $t \in \{0, 1, 2\}$.

[67], [68], [69], [70], [71], [72], [103], [104], [105] that learn a spatio-temporal model from training data, RED-PSM does not require ground-truth spatio-temporal training data, which is often not available or expensive to acquire. Thus RED-PSM offers the best of both worlds: a learned spatial model using readily available *static* training data, and unsupervised single-instance adaptation to the spatio-temporal measurement.

4) A novel and effective ADMM algorithm for the resulting new minimization problem enforces PSM as a hard constraint and incorporates the learned RED spatial regularizer.

5) The method is supported by theoretical analysis: we show convergence of the proposed PSM-based objective to a stationary point, and do so for a highly non-trivial learned regularizer. To date, the convergence of RED has been analyzed in the nonlinear measurement setting only for the phase retrieval problem [19]. Our analysis of RED-PSM is the first convergence result of RED in the bilinear (bi-convex), non-convex scenario. This convergence guarantee is of practical importance. First, it provides a mathematical justification to terminate the algorithm after sufficiently many iterations. Second, it eliminates concerns about the solution degrading after too many iterations as in the case of DIP-based methods.

6) Compared to a recent DIP-based [14] algorithm, RED-PSM achieves better reconstruction accuracy with orders of magnitude faster run times.

7) To improve and speed up the empirical convergence of RED-PSM, we use a particular fast projection-domain PSM initialization scheme [93], [94] for the spatial and temporal basis functions. The accelerated and reliable convergence with this initialization scheme is important for applications that require fast turn-around between acquisition and reconstruction.

8) A version of the approach with a patch-based regularizer is shown to provide almost equivalent reconstruction accuracy. This makes the proposed method conveniently scalable to high-resolution 3D or 4D settings.

An earlier and partial version of this work, missing, among other things, the detailed discussion of previous work, convergence analysis, and some of the experimental results, was presented at a conference [106].

II. PROBLEM STATEMENT

In a 2D setting, illustrated in Fig. 1, the goal in the dynamic tomography problem is to reconstruct a time-varying object $f(\mathbf{x}, t)$, $\mathbf{x} \in \mathbb{R}^2$ vanishing outside a disc of diameter D , from its projections

$$g(\cdot, \theta, t) = \mathcal{R}_\theta\{f(\mathbf{x}, t)\}$$

obtained using the Radon transform operator \mathcal{R}_θ at angle θ . Considering time-sequential sampling, in which only one projection is acquired at each time instant, and sampling uniform in time, the acquired measurements are

$$\{g(s, \theta_p, t_p)\}_{p=0}^{P-1}, \quad \forall s, t_p = p\Delta t, \quad (1)$$

where s is the offset of the line of integration from the origin (i.e., detector position), and P is the total number of projections (and temporal samples) acquired. The sampling of the s variable is assumed fine enough and is suppressed in the notation. The angular sampling scheme, the sequence $\{\theta_p\}_{p=0}^{P-1}$, with $\theta_p \in [0, 2\pi]$, is a free design parameter.

Our objective in dynamic tomography is to reconstruct the underlying dynamic object $\{f(\mathbf{x}, t_p)\}_{p=0}^{P-1}$ from the time-sequential projections in (1). The challenge is that because each projection belongs to a different object, the projections in (1) are inconsistent. Therefore, a conventional, e.g., filtered back-projection (FBP) reconstruction as for a static object results in significant reconstruction artifacts. This is to be expected, as the problem is severely ill-posed: an image with a D -pixels diameter requires more than D projections for artifact-free reconstruction, whereas only one projection is available per time-frame in the time-sequential acquisition (1). Several dynamic tomography methods [79], [89], [100] group temporally neighboring projections and assume the object is static during their acquisition. However, this reduces the temporal resolution, and any violation of this assumption (as in the case of time-sequential sampling) leads to mismodeling in the data fidelity terms.

III. PARTIALLY SEPARABLE MODELS (PSM)

For spatio-temporal inverse problems such as dynamic MRI and tomography, the underlying object can be accurately represented using a partially-separable model (PSM), which effectively introduces a factorized low-rank prior to the problem. For dynamic tomography, a PSM can represent the imaged object f , or its full set of projections g . In this paper we use an object-domain PSM.

The representation of a dynamic object $f(\mathbf{x}, t)$ by a K -th order partially separable model (PSM) is the series expansion

$$f(\mathbf{x}, t) = \sum_{k=0}^{K-1} \Lambda_k(\mathbf{x})\psi_k(t). \quad (2)$$

This model facilitates interpretability by separating the spatial structure from the temporal dynamics. Empirically, modest values of K provide high accuracy in applications to MR cardiac imaging [10], [107], [108]. Theoretical analysis [94] shows that for a spatially bandlimited object undergoing a time-varying affine transformation (i.e, combination of time-varying translation, scaling, and rotation) of bounded magnitude, a low order PSM provides a good approximation.

As detailed later, as a standard choice in modeling temporal functions, we use a d -dimensional representation with $d \ll P$. A similar idea was used in [11] to represent the low temporal bandwidth of cardiac motion in dMRI. Together with the PSM for the object this leads to a significant reduction in the number of representation parameters for a spatio-temporal object with a D -pixel diameter and P temporal samples: from $\approx PD^2$ to $\approx KD^2 + Kd$, with $K \ll P$. Because $d \ll D$, this corresponds

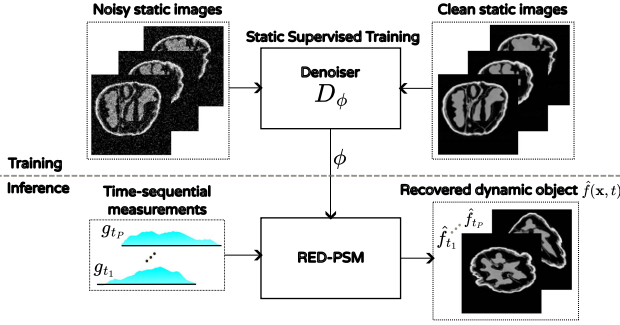


Fig. 2. The RED-PSM framework. The deep denoiser D_ϕ is trained on slices of *static* objects similar to the object of interest, and the learned spatial prior is used at inference time.

to a factor of $P/K \gg 1$ compression of the representation, providing an effective parsimonious model for the spatio-temporal object f .

Also, by propagating the PSM object model to the projection domain, it enables quantitative analysis [94] of the choice of optimal sequential projection angular schedule.

IV. PROPOSED METHOD: RED-PSM

The overall RED-PSM framework for recovering dynamic objects explained in this section is illustrated in Fig. 2.

A. Variational Formulation

We use a discretized version of the PSM (2) for the dynamic object, with the object $f(\cdot, t)$ at each time instant $t = 0, 1, \dots, P-1$, represented by a $N \times N$ -pixel image (a “time frame”, or “snapshot”). Vectorizing these images to vectors $f_t \in \mathbb{R}^{N^2}$, the entire dynamic object is the $N^2 \times P$ matrix $f = [f_0 \dots f_{P-1}]$. It will also be useful to extract individual time frames from f . Denoting the t -th column of the $P \times P$ identity matrix by e_t , we have $f e_t = f_t$, i.e., e_t extracts the t -th column of f . Applying the PSM model, we assume $f = \Lambda \Psi^T \in \mathbb{R}^{N^2 \times P}$, where the columns Λ_k and Ψ_k of $\Lambda \in \mathbb{R}^{N^2 \times K}$ and $\Psi \in \mathbb{R}^{P \times K}$ are the discretized spatial and temporal basis functions for the PSM representation, respectively.

Assuming that the x-ray detector has N bins, the projection of the object at time t is $g_t = g(\cdot, \theta_t, t) = R_{\theta_t} f_t \in \mathbb{R}^N$, where the measurement matrix $R_{\theta_t} \in \mathbb{R}^{N \times N^2}$ computes the projection at view angle θ_t .

We formulate the recovery of f as the solution $\hat{f} = \hat{\Lambda} \hat{\Psi}^T$ to the following variational problem

$$(\hat{\Lambda}, \hat{\Psi}) = \arg \min_{\Lambda, \Psi} \sum_{t=0}^{P-1} \left(\|R_{\theta_t} \Lambda \Psi^T e_t - g_t\|_2^2 + \lambda \rho(\Lambda \Psi^T e_t) \right) + \xi \|\Psi\|_F^2 + \xi \|\Lambda\|_F^2 \quad \text{where } \Psi = UZ. \quad (3)$$

The first term is the data fidelity term measuring the fit between available undersampled measurements g_t of the true object and the measurements obtained from the estimated object $f = \Lambda \Psi^T \in \mathbb{R}^{N^2 \times P}$. The second term with weight $\lambda > 0$ is a *spatial* regularizer injecting relevant spatial prior to the problem. It is applied to the PSM $\Lambda \Psi^T$ column by column, that is, to individual temporal image frames. The last two terms with

weight $\xi > 0$ prevent the bilinear factors from growing without limit.¹

Finally, the identity $\Psi = UZ$ is an implicit *temporal* regularizer that restricts the temporal basis functions Ψ to a d -dimensional subspace of \mathbb{R}^P spanned by a fixed basis $U \in \mathbb{R}^{P \times d}$ where $d \geq K$. The action of U on columns of Z may be interpreted as interpolation from d samples to the P -sample temporal basis functions. In practice, we incorporate this identity by explicit substitution (reparametrization of Ψ in terms of the free variable $Z \in \mathbb{Z}^{d \times K}$) into the objective, and the minimization in (3) over Ψ is thus replaced by minimization over Z . This reduces the number of degrees of freedom in Ψ to a fixed number dK , independent of the number P of temporal sampling instants. For notational conciseness, we do not display this constraint/reparametrization in the sequel, but it is used throughout.

B. Incorporating Regularization by Denoising

For the spatial regularizer $\rho(\cdot)$ we consider “Regularization by Denoising (RED)” [13]. RED proposes a recovery method using a denoising operator $D : \mathbb{R}^{N^2} \rightarrow \mathbb{R}^{N^2}$ in an explicit regularizer of the form

$$\rho(f_t) = \frac{1}{2} f_t^T (f_t - D(f_t)). \quad (4)$$

Recent studies using this regularizer provide impressive performances for various *static* image reconstruction tasks including high-dimensional cases [109] and phase retrieval [110]. This regularizer was also combined with a DIP-based fidelity term in [111]. However, we avoid this due to disadvantages related to speed, overfitting, and convergence guarantees.

While it provides significant flexibility for the type of denoisers that can be used, RED still requires D to be differentiable and locally homogeneous, and to satisfy the passivity condition $\|D(f_t)\| \leq \|f_t\|$, for the theoretical analysis of RED to apply.²

Next, we consider the optimization in (3). For the conventional variational formulation

$$\hat{f}_t = \arg \min_{f_t} \|R_{\theta_t} f_t - g_t\|_2^2 + \lambda \rho(f_t),$$

an efficient choice are iterative algorithms [13] that use the standard “early termination” approach [112], and only require a single use of the denoiser per iteration. However, the regularized PSM objective in (3) does not allow to propagate the RED updates on f to the respective basis functions efficiently. To overcome this difficulty, we perform a bilinear variable splitting $f = \Lambda \Psi^T$ and obtain our final formulation

$$\min_{f, \Lambda, \Psi} H(f, \Lambda, \Psi) \quad \text{s.t.} \quad f = \Lambda \Psi^T \quad (5)$$

$$\text{where } H(f, \Lambda, \Psi) = \sum_{t=0}^{P-1} \left(\|R_{\theta_t} \Lambda \Psi^T e_t - g_t\|_2^2 + \lambda \rho(f e_t) \right)$$

¹Inclusion of these Frobenius norm terms also happens to lead to connections with nuclear norm, but, as discussed in detail in Section V, our problem formulation (even if the RED regularizer is disregarded) is *not* equivalent to a formulation encouraging a low-rank solution by a nuclear norm or other Schatten norm penalty such as used in [44], [46], [47]. Instead, the factorized form $f = \Lambda \Psi^T$ (as in other PSM-based matrix recovery works such as [10], [11]) introduces this rank constraint explicitly.

²While many powerful denoisers have been demonstrated to satisfy these conditions in [13], recent work [20] provides another framework to explain the good performance of RED with denoisers not satisfying them.

Algorithm 1: RED-PSM.

input: $\Lambda^{(0)}, \Psi^{(0)}, \gamma^{(0)}, f^{(0)} = \Lambda^{(0)}\Psi^{(0)T}, \beta > 0, \lambda > 0, \xi > 0$
1: for $i \in \{1, \dots, I\}$ **do**
2: $\Lambda^{(i)} = \arg \min_{\Lambda} \{ \sum_t \|R_{\theta_t} \Lambda \Psi^{(i-1)T} e_t - g_t\|_2^2 + \frac{\beta}{2} \|\Lambda \Psi^{(i-1)T} - f^{(i-1)} + \gamma^{(i-1)}\|_F^2 + \xi \|\Lambda\|_F^2 \}$
3: $\Psi^{(i)} = \arg \min_{\Psi} \{ \sum_t \|R_{\theta_t} \Lambda^{(i)} \Psi^T e_t - g_t\|_2^2 + \frac{\beta}{2} \|\Lambda^{(i)} \Psi^T + \gamma^{(i-1)} - f^{(i-1)}\|_F^2 + \xi \|\Psi\|_F^2 \}$
4: $\forall t : f_t^{(i)} = \arg \min_{f_t} \{ \lambda \rho(f_t) + \frac{\beta}{2} \|(\Lambda^{(i)} \Psi^{(i)T} + \gamma^{(i-1)})e_t - f_t\|_2^2 \}$
5: $\gamma^{(i)} = \gamma^{(i-1)} + \Lambda^{(i)} \Psi^{(i)T} - f^{(i)}$
6: end for

Algorithm 2: RED-PSM With Efficient f step.

Notes: Inputs, and Lines 1–3 and 5–6 are the same as Algorithm 1. The f step is applied $\forall t$.
4: $\forall t : f_t^{(i)} = \frac{\lambda}{\lambda+\beta} D_{\phi}(f_t^{(i-1)}) + \frac{\beta}{\lambda+\beta} (\Lambda^{(i)} \Psi^{(i)T} + \gamma^{(i-1)})e_t$

$$+ \xi \|\Lambda\|_F^2 + \xi \|\Psi\|_F^2. \quad (6)$$

Since the PSM is enforced as a hard constraint, the estimated object f is constrained to have $\text{rank}(f) \leq K$. Problem (5) is non-convex even if ρ is convex, because of the product between unknowns Λ and Ψ .

We propose an algorithm based on ADMM to solve (5). To this end, we form the augmented Lagrangian in the scaled form, [113], [114]

$$\begin{aligned} \mathcal{L}_{\beta}[\Lambda, \Psi, f; \gamma] = & \sum_t \left(\left\| R_{\theta_t} \Lambda \Psi^T e_t - g_t \right\|_2^2 + \lambda \rho(f e_t) \right) \\ & + \xi \|\Psi\|_F^2 + \xi \|\Lambda\|_F^2 - \frac{\beta}{2} \|\gamma\|_F^2 + \frac{\beta}{2} \|\Lambda \Psi^T - f + \gamma\|_F^2, \end{aligned} \quad (7)$$

where $\gamma \in \mathbb{R}^{P \times N^2}$ represents the scaled dual variable associated with the constraint $f = \Lambda \Psi^T$ and $\beta > 0$ is the penalty parameter.

Then, ADMM can be used to solve (7) as in Algorithm 1.

Line 4 in Algorithm 1 then corresponds to the variational denoising for all t of the ‘‘pseudo image frame’’ $(\Lambda^{(i)} \Psi^{(i)T} + \gamma^{(i-1)})e_t$ with regularization $\lambda \rho(f_t)$. Instead of solving this denoising problem by running an iterative algorithm to convergence, we follow the RED approach [13], [20], and for each t replace the f_t update in Step 4 by a single fixed-point iteration step using the approach of early stopping [112] (Sec. 4.3.2), taking advantage of the gradient rule

$$\nabla \rho(f_t) = f_t - D_{\phi}(f_t) \quad (8)$$

where D_{ϕ} is the denoiser. This results in Algorithm 2, which requires only a single use of the denoiser per iteration of ADMM. Furthermore, as in Algorithm 1, the modified Line 4 in Algorithm 2 can be performed in parallel for all t .

C. Regularization Denoiser

The regularization denoiser D_{ϕ} has a deep neural network CNN (DnCNN) [115] architecture and is trained in a supervised manner on a training set of 2D slices $f_i \in \mathbb{R}^{N^2}, i = 1, \dots, N$

TABLE I

TIME AND SPACE COMPLEXITIES FOR TWO VARIANTS OF THE RED-PSM ALGORITHM FOR A SINGLE OUTER ITERATION OF THE BILINEAR ADMM

Complexity	PSM-fidelity	Patch-based
Space	$O(KN^2)$	$O(KN \frac{N_B^2}{s})$
Time	$O(KN^2 P M_i)$	$O(KN^2 P M_i (\frac{N_B}{s})^2)$

M_i is the number of inner iterations for each iteratively solved subproblem, and $N_B \ll N, s \leq N_B$.

of one or more *static* objects similar to the object of interest, assuming that such data will be available in the settings of interest. Thus, the RED steps are agnostic to the specific motion type. The training objective for the denoiser is

$$\min_{\phi} \sum_i \|f_i - D_{\phi}(\tilde{f}_i)\|_F^2 \quad \text{s.t.} \quad \tilde{f}_i = f_i + \eta_i, \quad \forall i, \quad (9)$$

where the injected noise $\eta_i \sim \mathcal{N}(0, \sigma_i^2 I)$ has noise level $\sigma_i \sim U[0, \sigma_{\max}]$ spanning a range of values, so that the denoiser learns to denoise data with various noise levels.

D. Computational Cost

Space (memory requirements) and time (operation count) complexities for two variants of the RED-PSM bilinear ADMM algorithm are shown in Table I. The operation counts are for a single outer iteration. As the number of outer iterations typically has a weak dependence on the size of the problem, the scaling shown tends to determine the run time. Furthermore, thanks to its structure, the algorithm also offers many opportunities for easy parallelization, so that actual runtime can be proportionally reduced by allocating greater computational resources. See the Supplementary Material Section VII-B for the detailed analysis of computational requirements.

The complexities for the proposed Algorithm 2 are given in the first column of Table I. Space complexity is dominated by the storage of Λ , and scales proportionally to image cross-section size in pixels N^2 , and the order K of the PSM. The time complexity is dominated by the computations of the gradient with respect to Λ , and scales proportionally to the size $N^2 P$ of the spatio-temporal object, PSM order K , and the number M_i of inner iterations used to solve optimization subproblems (Lines 2 and 3 in Algorithm 2).

Finally, in the second column of Table I, we investigate the patch-based version of the proposed algorithm described in Section VI-D6. Given a patch size $N_B \ll N$ and stride $s \leq N_B$, this alternative increases the operation count by $(\frac{N_B}{s})^2$ since it operates on overlapping patches. For example, for the settings in our experiments, $s = \frac{N_B}{2}$, increasing the operation count by a factor of 4. This is a modest price to pay, because this alternative also reduces the space complexity by a factor of $(\frac{N_B}{N})^2$. For example, for $N_B=8$ and $N=128$ as used in our experiments, this corresponds to a reduction by a factor of 144 in space complexity. Thus, this variant of RED-PSM enables scaling for high-dimensional and high-resolution settings.

V. CONVERGENCE ANALYSIS

In this section, we follow an approach similar to recent work on ADMM for a bilinear model [116] to analyze convergence. We show that under mild technical conditions, the objective in Algorithm 1 is guaranteed to converge (with increasing number

I of iterations) to a value corresponding to a stationary point of the Lagrangian, that is, satisfying the necessary conditions for first order optimality.

In practice, Algorithm 2 with the efficient f step version, which we implemented and used in the experiments reported in Section VI, has better run times, and rapid empirical convergence. However, its analysis requires additional steps, which are not particularly illuminating. Therefore we focus on the analysis of the nominal Algorithm 1.

In spite of the similarity, at a high level, to the problem and analysis in [116], our problem formulation and algorithm differ in several aspects, which require modifications in the analysis and proof. In particular, different to [116], where the bilinear form only appears in the constraint, the bilinear form appears both in our objective function and constraint. To satisfy strong convexity for the respective subproblems, our objective also uses the Frobenius norm terms for the PSM factors, instead of the interim proximal terms in [116]. Moreover, to allow the efficient f -step in Algorithm 2, we use a different order of updates than [116]. Importantly, we account for the explicit RED regularization and its required properties in the proof of convergence. Our analysis includes a proof (similar to [117]) of the boundedness of the iterates to justify the existence of points of accumulation of the iterate sequence, which appears to have been inadvertently left out in [116]. However, different to [117], which has objectives with linear constraints, our objective (3) has a bilinear constraint.

To simplify the notation, we replace the separate computation of the projections of time t -frames in the data fidelity term by using the operator $\bar{R} : \mathbb{R}^{N^2 \times P} \rightarrow \mathbb{R}^{N \times P}$ that computes the entire set of P projections at view angles θ_t of the image frames at times t , $t = 1, \dots, P$ of dynamic image f , i.e., of each of its columns indexed by t , producing $g = \bar{R}f \in \mathbb{R}^{N \times P}$. When applied to the PSM, \bar{R} performs $R_{\theta_t} \Lambda \Psi^T e_t$ for each t . We also aggregate the contribution of the RED regularizer into $\bar{\rho}(f) \triangleq \sum_{t=0}^{P-1} \rho(f_t) : \mathbb{R}^{N^2 \times P} \rightarrow \mathbb{R}$, and the denoiser into $\bar{D} : \mathbb{R}^{N^2 \times P} \rightarrow \mathbb{R}^{N^2 \times P}$, which performs D for each column of f indexed by t . Then, the Lagrangian function with dual variable γ can be rewritten as

$$\begin{aligned} \mathcal{L}[\Lambda, \Psi, f; \gamma] &= \|\bar{R}\Lambda\Psi^T - g\|_F^2 + \lambda\bar{\rho}(f) \\ &\quad + \xi\|\Lambda\|_F^2 + \xi\|\Psi\|_F^2 + \beta\langle\gamma, (\Lambda\Psi^T - f)\rangle, \end{aligned} \quad (10)$$

where the inner product is defined as $\langle A, B \rangle = \text{Tr}(A^T B)$.

The corresponding augmented Lagrangian is

$$\begin{aligned} \mathcal{L}_\beta[\Lambda, \Psi, f; \gamma] &= \|\bar{R}\Lambda\Psi^T - g\|_F^2 + \lambda\bar{\rho}(f) \\ &\quad + \xi\|\Lambda\|_F^2 + \xi\|\Psi\|_F^2 + \beta\langle\gamma, (\Lambda\Psi^T - f)\rangle + \frac{\beta}{2}\|\Lambda\Psi^T - f\|_F^2. \end{aligned} \quad (11)$$

Then, we can state the subproblems with respect to each primal variable in Algorithm 1 as

$$\begin{aligned} S_\Lambda &= \|\bar{R}\Lambda\Psi^{(i-1)T} - g\|_F^2 + \xi\|\Lambda\|_F^2 \\ &\quad + \frac{\beta}{2}\|\Lambda\Psi^{(i-1)T} - f^{(i-1)} + \gamma^{(i-1)}\|_F^2 \end{aligned} \quad (12)$$

$$\begin{aligned} S_\Psi &= \|\bar{R}\Lambda^{(i)}\Psi^T - g\|_F^2 + \xi\|\Psi\|_F^2 \\ &\quad + \frac{\beta}{2}\|\Lambda^{(i)}\Psi^T - f^{(i-1)} + \gamma^{(i-1)}\|_F^2, \text{ and} \end{aligned} \quad (13)$$

$$S_f = \lambda\bar{\rho}(f) + \frac{\beta}{2}\|\Lambda^{(i)}\Psi^{(i)T} - f + \gamma^{(i-1)}\|_F^2. \quad (14)$$

Algorithm 1 will be shown in Theorem 1 to converge to a stationary point of Problem (5), which is defined below.

Definition 1: (Stationary solution for Problem (5)). The point $W^* = (\Lambda^*, \Psi^*, f^*, \gamma^*)$ is a stationary solution of the problem (5) if it satisfies the stationarity and primal feasibility conditions for the variables of the Lagrangian in (10):

$$\nabla_f \mathcal{L}(\Lambda^*, \Psi^*, f^*; \gamma^*) = \lambda \nabla \bar{\rho}(f^*) - \beta \gamma^* = 0; \quad (15a)$$

$$\begin{aligned} \nabla_\Lambda \mathcal{L}(\Lambda^*, \Psi^*, f^*; \gamma^*) &= (2\bar{R}^T(\bar{R}\Lambda^*\Psi^{*T} - g) + \beta\gamma^*)\Psi^* \\ &\quad + 2\xi\Lambda^* = 0; \end{aligned} \quad (15b)$$

$$\begin{aligned} \nabla_\Psi \mathcal{L}(\Lambda^*, \Psi^*, f^*; \gamma^*) &= (2\bar{R}^T(\bar{R}\Lambda^*\Psi^{*T} - g) + \beta\gamma^*)^T \Lambda^* \\ &\quad + 2\xi\Psi^* = 0; \end{aligned} \quad (15c)$$

$$f^* = (\Lambda\Psi^T)^{(*)}. \quad (15d)$$

By its definition, a stationary solution of (5) satisfies the necessary first-order conditions for optimality.

To state Theorem 1, we need to introduce some additional definitions pertaining to the denoiser D used in RED.

Definition 2: (Strong Passivity [13]). A function D is strongly passive if $\|D(f)\| \leq \|f\|$.

Moreover, we assume the gradient rule of RED (8). This is shown to hold true when the denoiser D is assumed locally homogeneous (LH) [13] and Jacobian symmetric (JS) [20]. It was shown [13] that several popular denoisers are practically LH, and that for (8) to hold for the explicit regularizer ρ in (4), JS is necessary [20].

In our analysis, we only require the gradient rule (8) to hold, rather than the explicit form (4) of the regularizer. Even for denoisers without LH and JS, the gradient rule can be used to define the fixed point iteration in the algorithm [20], and, more recent works on using denoisers for regularization [118], [119], [120], [121], [122] (in particular [120]) shows how to learn a denoiser jointly with the gradient so that the gradient rule is satisfied.

Finally, we will also assume Lipschitz continuity of the denoiser with some Lipschitz constant L_D , i.e., $\|D_\phi(f_1) - D_\phi(f_2)\|_2 \leq L_D\|f_1 - f_2\|_2$ for all images $f_1, f_2 \in \mathbb{R}^{N^2}$. This means that the denoiser has some finite gain L_D , which is satisfied by any reasonable denoiser.

Our main convergence result is the following.

Theorem 1: Suppose that the denoiser D_ϕ satisfies the gradient rule in (8), and is Lipschitz continuous with some Lipschitz constant L_D , and strongly passive.

Then, if $\beta > 2L$ where $L \triangleq \lambda(1 + L_D)$, Algorithm 1 converges globally (i.e., regardless of initialization) to a stationary solution (f^*, Λ^*, Ψ^*) of (5) in the following sense:

- i) The sequence of values of the objective H converges to a limit $H^* = H(f^*, \Lambda^*, \Psi^*)$; and
- ii) The iterates generated by Algorithm 1 converge subsequentially to a stationary solution $(f^*, \Lambda^*, \Psi^*, \gamma^*)$ of (5), that is, any accumulation point (of which there is at least one) of the sequence $(f^{(i)}, \Lambda^{(i)}, \Psi^{(i)}, \gamma^{(i)})$ is a stationary solution of (5).

The proof of Theorem 1 is provided in the supplementary material Section D.

Regarding convergence to the globally optimal solution, more can be said. Problem (3) is non-convex even if $\bar{\rho}$ is convex, because of the product between unknowns Λ and Ψ . However, similar to Theorem 7.1 of [123] (which generalizes arguments in [124], [125]), thanks to the inclusion of the Frobenius norms of the factors Λ and Ψ in the bilinear form, when $\bar{\rho}$ is convex (i.e., when the gradient rule (8) is assumed to hold), the global minimum $\hat{f} = \hat{\Lambda}\hat{\Psi}^T$ of (3) can be shown to coincide with the global minimum in

$$\begin{aligned} \hat{f} = \arg \min_f \sum_{t=0}^{P-1} \|R_{\theta_t} f_t - g_t\|_2^2 + \lambda \bar{\rho}(f) + 2\xi \|f\|_* \\ \text{s.t. rank}(f) \leq K, \end{aligned} \quad (16)$$

where the penalty weight $\xi > 0$ is identical to the Frobenius norm penalty in our RED-PSM objective (5).

Importantly, Problem (16) is non-convex too, because of the rank constraint. In other words, it is *not* equivalent to a conventional nuclear norm penalized problem. Hence, although the Frobenius norm penalties on the PSM factors in our problem formulation (3) lead to a connection to nuclear norm, Problem (3) is *not equivalent* to a formulation encouraging a low-rank solution by a nuclear norm penalty.

However, returning to the question of determining global optimality of a candidate solution to (16), consider the convex version of Problem (16), without the rank constraint, and denote its solution by \hat{f}^{Convex} . If it so happens that $\text{rank}(\hat{f}^{\text{Convex}}) \leq K$, then the low-rank constraint in Problem (16) is not active, and $\hat{f}^{\text{Convex}} = \hat{f}$. It then follows [123] that the optimality conditions of the convex problem (without the rank constraint) can be used to assess the global optimality of a candidate solution $\hat{f} = \hat{\Lambda}\hat{\Psi}^T$ produced by the algorithm solving (3).

VI. EXPERIMENTS

A. Datasets

Three categories of data sets are used in this work.

Walnut Dataset: We use the CT reconstructions of two different (static) walnut objects from the publicly available 3D walnut CT dataset [126]. We create a dynamic test object by synthetically warping the central axial slice of one of the walnut objects using a sinusoidal piecewise-affine time-varying warp [127]. To be precise, the image is divided into a $N \times N$ uniformly spaced rectangular grid, and the following vertical displacement is applied on each row separately to drive the temporally varying warp $\Delta_{n,t} = -C(t) \sin(3\pi n/N)$, $n \in \{0, \dots, N-1\}$, where $C(t)$ is a linearly increasing function of t and $C(0) = 0$. *Static* axial, coronal, and sagittal slices of the other walnut object are used to train the denoiser D_ϕ .

Compressed Object Dataset: The compressed object data set is obtained from a materials science experiment [128] with a sequence of nine increasing compression (loading) steps applied to an object, with a full set of radiographic projections collected (using Carl Zeiss Xradia 520 Versa) and reconstructed by the instrument’s software at each step. Using this quasi-static data set, a fixed axial slice is extracted from each of the 9 reconstructions. These nine extracted slices corresponding to nine time points are interpolated to P time frames using a recent deep learning-based video interpolation algorithm [129]. The denoiser D_ϕ for our experiments on this data set was trained using the axial slices




Fig. 3. Ground-truth frames uniformly sampled in time for $P = 4$, for the time-varying walnut (top) and compressed object (bottom).

of the *static* pre-compression and post-compression versions of the object, which would be available in actual dynamic loading experiments of this type.

In a conference submission [130] citing this work, the method has been applied to additional tomographic scenarios of materials compressed under load, further confirming the advantage of the proposed method and algorithm over the compared methods.

We note that all algorithms compared in this paper are agnostic to both the synthetic warp applied on the static walnut slice and to the data-driven interpolation method used for the compressed object.

Spatio-temporal projection data for each dataset is simulated by a parallel-beam projection with $N=128$ detector bins of the dynamic phantoms, a single projection at each of the P time instants. The sequence of projection angles $\{\theta_t\}_{t=0}^{P-1}$ (a free experimental design parameter) was chosen to follow the bit-reversed view angle sampling scheme, which has been shown [93] to provide an especially favorable conditioning of the recovery problem. The simulated measurements are corrupted using AWGN with standard deviation $\sigma = 5 \cdot 10^{-3}$. This noise level leads to the FBP (with Ram-Lak filter) of the full set of $P=512$ projections at each time instant having a PSNR of approximately 46 dB. When, in the actual experiments with sequentially sampled data, only $1/P$ of this data is used, the PSNR of the reconstruction may be expected to be lower.

Ground-truth frames for $P = 4$ are shown in Fig. 3.

Cardiac dMRI Dataset: For a more direct comparison with the setting and data used in dMRI works, we also test RED-PSM on the “retrospective” cardiac dMRI data in [14]. The data includes 23 distinct time frames for one cardiac cycle. Details of the data and experiments are in Section VI-D4.

B. Comparison Benchmarks

PSM-TV: Similar to the proposed approach, this algorithm also uses a PSM to represent the object, but instead of the RED regularizer, the regularization penalizes the discrete 2D total variation of the temporal frames of f at each time instant. To this end, the constraint $f = \Lambda\Psi^T$ is implemented by substitution into the objective in (5), and the definition of ρ is changed to $\rho(\cdot) = \text{TV}(\cdot)$, and the rest of the objective is kept the same. We consider spatial (PSM-TV-S) and spatiotemporal (PSM-TV-ST) alternatives of TV. The former computes TV only spatially in a single frame at each t whereas the latter also computes differences between temporally adjacent pixels at $t-1$ and $t+1$. The unconstrained problem is then solved for $\{\hat{\Lambda}, \hat{\Psi}\}$ (using Adam optimizer in Pytorch). Finally, the estimated object is obtained as $\hat{f} = \hat{\Lambda}\hat{\Psi}^T$.

TD-DIP [14]: TD-DIP is a recent method based on the Deep Image Prior (DIP) approach.³ It uses a mapping network M_α and a generative model N_β in cascade to obtain the estimated object at each time instant f_t from fixed and handcrafted latent representations χ_t . Because TD-DIP was originally proposed for dynamic MRI, we modified the objective minimally for dynamic tomography as

$$\min_{\alpha, \beta} \sum_t \|g_t - R_{\theta(t)}((N_\beta \circ M_\alpha)(\chi_t))\|^2. \quad (17)$$

For the comparisons in this work, the mapping network and generator architectures, latent representation dimensionality, optimizer, learning rate, and decay schemes are identical to those in the available implementation [131]. The original work focuses on the beating heart problem and thus proposes a helix-shaped manifold for χ_t with number of cycles equal to the number of heartbeats during measurement acquisition. Since we do not have a repetition assumption for the motions of the walnut and compressed object, we use a linear manifold as explained in the original paper [14]. Thus, for clarity, in Section VI-D the method is sometimes denoted as “TD-DIP (L)”.

C. Experimental Settings

All methods are run on a workstation with an Intel(R) Xeon(R) Gold 5320 CPU and NVIDIA RTX A6000 GPU. In practice, we used a minor variation of Algorithm 2, where we combined the subproblems for Λ and Ψ , and minimized with respect to both basis functions simultaneously using gradient descent with Adam [132] optimizer.

Denoiser and training: Each convolutional layer in the denoiser network is followed by a ReLU nonlinearity except for the final single-channel output layer. We tested both direct and residual DnCNN denoisers, where the former predicts the denoised image and the latter estimates the noise from the input. We use the denoiser type that performs better for each object, but the differences are minor. Further architectural details for the denoisers in our experiments are in Table VII in the Supplementary Material. We use three pre-trained denoisers, one for each of the three object types. In each case, the same pre-trained denoiser was used for all values of P .

The upper limit for noise level used in training the denoiser was set to $\sigma_{\max} = 5 \cdot 10^{-2}$. For the dynamic walnut object, the denoiser D_ϕ is trained on the central 200 axial, 200 sagittal, and 200 coronal slices of another static walnut CT reconstruction downsampled to size 128×128 . For the compressed object, axial slices of pre-compression and post-compression *static* versions of the object, containing 462 slices in total, are used to train D_ϕ . For the cardiac MRI setting, the denoiser was trained on the *static* MRI training slices of the ACDC dataset [133]. For all datasets, D_ϕ is trained for 500 epochs using the Adam optimizer with a learning rate of $5 \cdot 10^{-3}$. As mentioned above, we evaluate both direct and residual DnCNN denoisers.

Temporal Basis: In compressed material and cardiac dMRI data experiments, we use the parametrization $\Psi = UZ$ with a

³It would be interesting to include yet another comparison benchmark (also developed for dMRI) – the DIP-based PSM approach [74]. However, as an implementation of this method was unavailable, and due to potential issues with replicating its performance and adapting to our CT problem (specific initialization and framework that use other MRI algorithms), we were unable to do so.

fixed basis U that corresponds to a cubic spline interpolator, and for the warped walnut we use DCT-II, to interpolate the low-dimensional temporal representation Z to Ψ .

Initialization: Unless stated otherwise, the spatial and temporal basis functions are initialized using the SVD truncated to the rank of the dynamic object estimate produced by a recent projection-domain PSM-based method “ProSep” [94]. If the ProSep estimate has rank smaller than K , the remaining basis functions are initialized as 0. Otherwise, all spatial basis functions are initialized as 0 and the temporal latent representations z_k are initialized randomly as $z_k \sim \mathcal{N}(0, I)$.

Tomographic Acquisition Scheme: All methods mentioned in this paper use the bit-reversed angular sampling scheme, over the range $[0, \pi]$. For time-sequential acquisition, the bit-reversed scheme was shown [93], [94] to provide favorable results via better conditioning of the forward model in comparison to alternatives. In a standard CT scanner, the speed of rotation might have to be significantly increased to implement this scheme, possibly leading to greater motion blur. However, for several scenarios this would not be much of an issue. These include radial acquisition in MRI; a CT scanner with electronic beam deflection [134]; and settings where the acquisition time is dominated by the time to acquire each view rather than the rotation time, e.g., micro-CT, or imaging of dense materials.

To help address the challenge in implementing the bit-reversed scheme in other, more physically constrained settings, the number of distinct view angles \hat{P} can be reduced and these views can be repeated periodically without a performance drop as also shown in Fig. 11. With a reduced \hat{P} , the bit-reversed scheme can be implemented more conveniently, e.g., by multiple source-detector pairs, or by carbon nanotube sources [135], [136].

Run Times: For $P = 256$ and using the specified computational resources and parameter settings, to reach the peak PSNR during optimization, RED-PSM with ProSep initialization requires $50 < \text{iterations} < 150$ taking about 2 to 6 minutes whereas TD-DIP with batch size P typically requires > 30 k steps, taking about 3.5 hours to complete. Hence, RED-PSM provides a speedup over TD-DIP by a factor of 35 to 105. Depending on the parameter configuration, the speedup factor may vary. However, the proposed method provides a significant run time reduction in all cases.

Evaluation Metrics: Four quantitative metrics were implemented for comparing different method performances: i) the *peak signal-to-noise ratio* (PSNR) in dB; ii) the *structural similarity index* (SSIM) [137]; iii) the *mean absolute error* (MAE); and iv) the *high-frequency error norm* (HFEN) [138]. The latter is defined as $HFEN(f, f_r) = \|\text{LoG}(f) - \text{LoG}(f_r)\|_2$ where LoG is a rotationally symmetric Laplacian of Gaussian filter with a standard deviation of 1.5 pixels.

D. Results

1) *Reconstruction Accuracies for Different P*: RED-PSM is compared in Fig. 4 and Table II with PSM-TV-S, PSM-TV-ST and TD-DIP (using, for both a-periodic objects, a linear latent representation). Parameter configurations for the experiments are provided in Table V in Supplementary Material A. While Fig. 4 facilitates the assessment and comparison of trends in the metrics for different P for the various methods, Table II provides the detail for a more precise quantitative comparison. In fact, since the plotted range of metrics such as PSNR in Fig. 4




Fig. 4. Reconstruction metrics for the time-varying walnut and compressed material vs. P using different methods. For TD-DIP, the metrics reported are assuming a “stopping oracle” that stops the iterations at the best PSNR reconstruction. The minimum and maximum values over three different runs with different random initial values are shown with bars, and the mean values are connected by dashed lines.

TABLE II
RECONSTRUCTION ACCURACIES FOR FOR DIFFERENT P

P	Method	(a) Walnut				(b) Compressed Material			
		PSNR (dB)	SSIM	MAE (1e-3)	HFEN	PSNR (dB)	SSIM	MAE (1e-2)	HFEN
32	PSM-TV-S (R)	21.1	0.678	2.9	1.04	22.5	0.697	2.9	11.6
	PSM-TV-ST (R)	21.3	0.731	2.6	0.99	24.1	0.702	2.5	9.3
	TD-DIP (L)	22.5	0.882	2.1	0.87	25.4	0.909	1.5	8.1
	RED-PSM (Pr)	22.8	0.911	1.5	0.78	23.9	0.916	1.2	9.9
64	PSM-TV-S (R)	24.0	0.879	1.9	0.94	25.6	0.845	1.8	10.7
	PSM-TV-ST (R)	24.6	0.897	1.7	0.83	27.7	0.874	1.3	8.2
	TD-DIP (L)	25.6	0.916	1.5	0.74	29.0	0.947	1.1	6.5
	RED-PSM (Pr)	26.4	0.958	0.9	0.57	29.3	0.952	0.5	6.0
128	PSM-TV-S (R)	26.8	0.912	1.4	0.78	29.2	0.907	1.1	9.6
	PSM-TV-ST (R)	27.7	0.931	1.2	0.63	31.2	0.942	0.9	7.0
	TD-DIP (L)	29.0	0.953	1.0	0.50	32.4	0.953	0.8	5.1
	RED-PSM (Pr)	30.3	0.979	0.6	0.40	33.1	0.972	0.4	4.7
256	PSM-TV-S (R)	30.1	0.934	1.0	0.63	31.9	0.963	0.7	8.0
	PSM-TV-ST (R)	31.4	0.942	0.8	0.45	32.7	0.978	0.6	6.8
	TD-DIP (L)	31.7	0.966	0.8	0.45	34.7	0.966	0.7	4.6
	RED-PSM (Pr)	34.2	0.989	0.4	0.32	35.9	0.986	0.3	4.4

Random (R) and ProSep (Pr) initialization for the PSM methods. For TDDIP, the reported accuracies are for the best PSNR using a “stopping oracle”, averaged over three runs with random initial conditions.

is very large due to varying P , Table II emphasizes important differences between methods for the same P .

Remarks:

1) The scale differences in the MAE and HFEN between the two objects are due to working with un-normalized densities.

2) In all TD-DIP experiments, optimization is stopped early to achieve the best PSNR reconstructions, assuming such a “stopping oracle” is available. In a more realistic setting, in the absence of this oracle, TD-DIP optimization with continued iterations suffers from overfitting, and produces in these experiments degraded results. For instance, for the warped walnut slice with 40 k iterations, PSNR, SSIM, and MAE degrade significantly to 19.5 dB, 0.822, and $3.1 \cdot 10^{-3}$ for $P = 32$, and 24.1 dB, 0.824, and $2.1 \cdot 10^{-3}$ for $P = 64$. Thanks to the accurate spatial prior and the convergence properties, we do not encounter such a problem for RED-PSM.

3) Because the performance of TD-DIP varies with initialization, in each experiment we ran it three times, with different

random initialization, each time using a “stopping oracle” to obtain the best PSNR for the given initialization. Table II reports the average of the best PSNR reconstruction accuracies for TD-DIP in these three runs. Fig. 4, complements this information, by showing in addition to the average results, also the best and worst of these runs (still using the stopping oracle to get the best PSNR per run).

As expected, for all methods, the estimates improve with increasing P . In terms of PSNR, the proposed algorithm performs on par with or usually better than TD-DIP (for all but the lowest P on the compressed material object), and consistently better than PSM-TV-S and PSM-TV-ST. The PSNR improvement of RED-PSM over TD-DIP enhances with increasing P , reaching 2.4 dB for the time-varying walnut with $P = 256$. Moreover, for the other three metrics, SSIM, MAE, and HFEN, the improvement of RED-PSM over other algorithms is more significant. Specifically, the reduction in MAE reaches up to and exceeds %50 for both objects for $P = 256$. These observations are valid for both objects, however, RED-PSM provides slightly greater improvement over TD-DIP in the walnut case.

We would like to emphasize the significance of these reconstruction quality improvements by comparing them to some representative examples. TD-DIP reports up to 3 dB PSNR and 0.005 SSIM improvement with respect to an older dMRI method in a *single* scenario. While providing comparable PSNR improvements (2.4 dB), we are able to provide four times larger SSIM (0.02) improvements *relative to TD-DIP itself*. Other recent dMRI [73] and dynamic photo-acoustic tomography [102] methods improve 1.5 dB, and at most 1 dB and often no improvement over their respective benchmarks, again in single imaging scenario for each method.

Also, based on the run-time improvement ($\sim 3x$) of [73] with respect to TD-DIP, RED-PSM is still much ($\sim 10x-30x$) faster, than [73], and more interpretable.

Fig. 5 compares the reconstructions for both objects at two different values of t for $P = 256$. As expected, PSM-TV-S performs the worst among the compared methods and provides, for both objects, blurry reconstructions lacking finer details. The TD-DIP reconstructions improve somewhat over PSM-TV-S and PSM-TV-ST, but contain visible noise-like artifacts on the

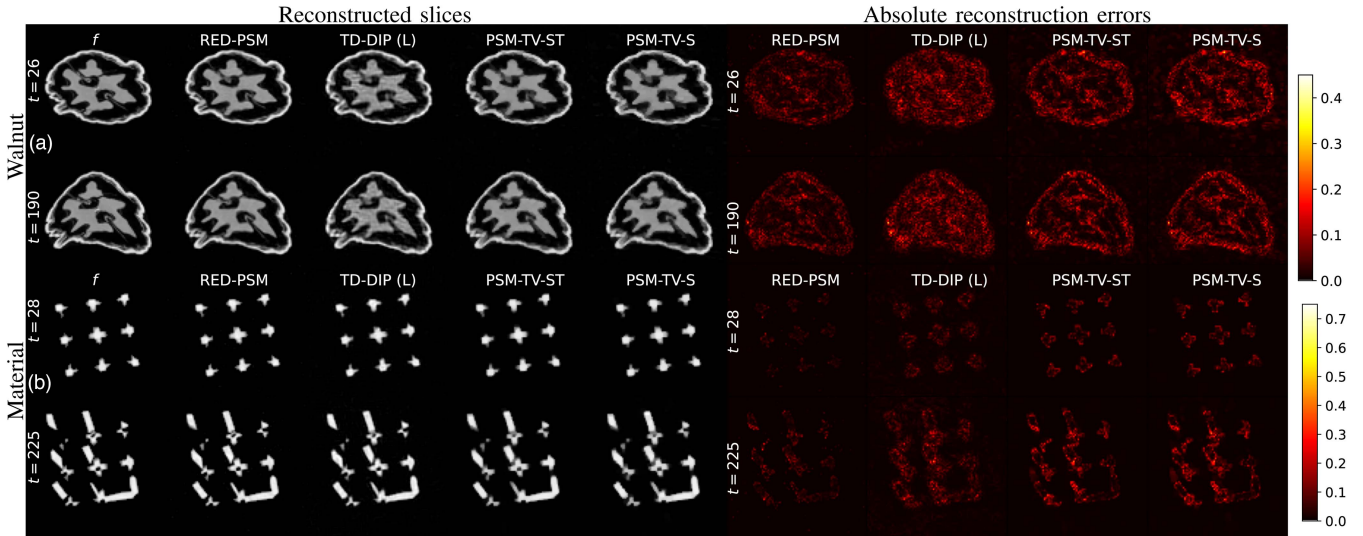


Fig. 5. Comparison of reconstructed object frames at two time instants using different methods for $P = 256$, and the corresponding normalized absolute reconstruction errors for (a) the time-varying walnut, and (b) compressed object.




Fig. 6. Comparison of reconstructed $x-t$ slices (top) and corresponding normalized absolute error (bottom) using different methods for $P = 256$. The cross-section location is indicated on the static $t = 0$ object with a yellow line. The $x-y-t$ coordinates are indicated in white text on the static object and bottom left absolute error figure.

piecewise constant regions of the walnut object, which are alleviated by RED-PSM. This is manifested also in the absolute difference figures, with error for TD-DIP distributed throughout the interior regions of the walnut. Also, around the shell of the walnut, RED-PSM is further able to preserve sharper details. For the compressed material, in comparison to TD-DIP, RED-PSM shows reduced absolute error almost uniformly over the object. This difference is more prominent around the highly dynamic regions of the object, emphasizing the advantage of the proposed method. For a better understanding of RED-PSM results, we also display the reconstructed spatial and temporal basis Λ and Ψ for the time-varying walnut scenario in Supplementary Material C.

Reconstructed $x-t$ “slices” through the dynamic walnut are compared in Fig. 6. The location of the $x-t$ slice is highlighted on the $t = 0$ static $x-y$ frame by a yellow line. Consistent with the comparison in Fig. 5, RED-PSM provides reduced absolute error values throughout the respective $x-t$ slice. Also, as more apparent on the error figures, TD-DIP leads to higher background errors.

Finally, a zoomed-in comparison of the time-varying walnut object for another time instant for $P = 256$ is provided in Fig. 7. The comparison shows the better performance of RED-PSM at recovering the finer details clearly.

2) *PSNR Vs. t Comparisons*: To complement the cumulative metrics in Fig. 4 and Table II and the “snapshot” qualitative comparisons in Fig. 5, we study how the reconstructed




Fig. 7. Zoomed-in walnut reconstruction comparison for $P = 256$.

image frame PSNRs vary over the reconstructed time interval. The per frame PSNRs (in dB) of the walnut and compressed object reconstructed with $P = 256$ by the different methods are shown in Fig. 8 as a function of t . For TD-DIP, the best PSNR obtained using a “stopping oracle” is reported, with the red shading indicating, for each t , the interval between the highest and lowest PSNR in three runs with different random initial conditions. For the warped walnut object, RED-PSM provides consistently better PSNR than the best-case TD-DIP for all t . For the compressed object, the same is true at about 70% of t points. Fig. 8 also shows transient effects at the beginning and the end for both objects and all methods. In scenarios such as the object compression experiment, in which the initial and final state are static and could be measured using multiple projections, such transients could be eliminated. Similarly, in quasi-periodic




Fig. 8. Reconstruction PSNR vs. t for the (a) time-varying walnut, and (b) compressed material for $P = 256$. The red shading for TD-DIP indicates, for each t , the interval between the best and worst PSNR with the early stopping oracle explained in Section VI-D1 in three runs with different random initial conditions.




Fig. 9. Advantage of ProSep-based vs. random initialization of RED-PSM (see Sec. VI-C). For random initialization, the area in blue between the best and the worst PSNR for each iteration highlights the varying performances of five different runs with different random initializations.

scenarios such as cardiac imaging the effect of such transients would be minimal.

3) *Effect of Initialization*: The initialization of Λ , Ψ , and f plays an important role in the performance and convergence speed of RED-PSM. We observe significant speed-up when rather than a random initialization, we initialize the algorithm with ProSep [93] estimated reconstruction. Fig. 9 shows PSNR vs. iterations comparison for different initialization techniques for the dynamic walnut object with $P = 256$. The rest of the parameters were selected identical to those indicated in Supplementary Material Table V. This experiment highlights the advantages of initializing with ProSep estimated basis functions: eliminating the need for multiple runs for a best-case result; and speeding up convergence considerably.

Combined also with the convergent algorithm eliminating the need for an unrealistic stopping oracle and the theoretical analysis, RED-PSM provides improved reliability which is of practical significance.

4) *Cardiac dMRI Data Experiments*: In this setting, different to previous experiments, we used 4 k -space radial lines (“spokes”) per frame at the bit-reversed angles. We used 1.4 and 2.8 cardiac cycles, with $23 \times 4 = 92$ spokes/cycle, for a total of $P=128$ and $P=256$ spokes. The problem is still severely under-sampled compared to the experiment in [14] where 13 spokes are used per frame for 13 cycles, for a total of $13 \times 23 \times 13 = 3,887$ spokes.

Since the data is periodic, we also tested the helix latent scheme (H) for TD-DIP.

The metrics in Table III and the qualitative comparison in Fig. 10 with zoomed-in reconstructions and absolute error maps, show that RED-PSM performs better than both versions of TD-DIP.

In [73], the authors compare their method to an *older* version of TD-DIP without the improved latent representation prior scheme for cardiac dMRI and report 1.5 dB PSNR improvement.

TABLE III
RECONSTRUCTION ACCURACIES FOR RED-PSM AND TD-DIP FOR THE RETROSPECTIVE dMRI DATA [14]

P	Method	PSNR (dB)	SSIM	MAE (1e-2)	HFEN
128	TD-DIP (L)	34.6	0.923	1.7	3.38
	TD-DIP (H)	34.6	0.928	1.7	3.32
	RED-PSM	36.6	0.939	1.4	2.73
256	TD-DIP (L)	36.2	0.948	1.4	3.47
	TD-DIP (H)	36.4	0.947	1.4	3.44
	RED-PSM	38.4	0.962	1.1	2.99




Fig. 10. Reconstructed frames for $P = 256$ for retrospective dMRI data [14] with zoomed-in frames (middle row), and absolute reconstruction errors (last row).




Fig. 11. Reconstruction PSNR and SSIM for the time-varying walnut vs. the number of distinct view angles \hat{P} using RED-PSM with different total number of views P .

In RED-PSM, we exceed this improvement on the latest version of TD-DIP in multiple scenarios.

5) *Acquisition With Smaller Number of Distinct View Angles*: Since obtaining time-sequential projections from different angles in a sufficiently short time period can be physically challenging, we also test the performance of RED-PSM with an acquisition scheme that may be easier to realize physically: keeping the same total number of P projections, but taken at a smaller number \hat{P} (called “period”) of *distinct* view angles, which are also obtained using the bit-reversed angular sampling scheme. The comparison in Fig. 11 shows that it is possible to use up to 1/8-th of the distinct view angles without performance loss for RED-PSM.

6) *Patch-Based RED Denoiser*: To improve the scalability of the method to higher resolution and/or 3D dynamic object, conveniently, the objective in (5) can be manipulated to operate

TABLE IV
PERFORMANCE COMPARISON FOR DIFFERENT DENOISER TYPES FOR RED-PSM

Object	P	Denoiser	PSNR(dB)	SSIM	MAE(1e-3)	HFEN
Walnut	256	Full-image	33.8	0.987	0.4	0.32
		Patch-based	33.7	0.989	0.4	0.35
Material	256	Full-image	35.9	0.986	2.5	4.43
		Patch-based	35.6	0.981	2.7	4.56

Patch-based denoisers also use DnCNN and have the same configuration and training policy as the full image denoiser.

on the patches of temporal image frames of the time-varying object. This circumvents the need to store the complete image frame at a given time, and also enables the denoiser D_ϕ to be both trained and to operate on patches of image frames. To showcase the potential of the suggested scheme, we replace the full-size D_ϕ with a patch-based counterpart in the RED step and compare the performance with the originally proposed method for 2D dynamic objects.

The patch-based denoiser for RED updates is trained using

$$\min_{\phi} \sum_i \sum_l \|B_l f_i - D_\phi(B_l \tilde{f}_i)\|_F^2 \text{ s.t. } \tilde{f}_i = f_i + \eta_i, \forall i,$$

where η_i, σ_i are set as in (9), and B_l , with $l \in \{0, \dots, L-1\}$, is the operator to extract the l -th patch of the image. The denoiser D_ϕ operates separately on each patch.

To train the patch-based D_ϕ , uniformly random rotations (multiples of $\frac{\pi}{2}$), and random horizontal/vertical flips, each with $\frac{1}{2}$ probability, were used for data augmentation. The patch size was chosen 8×8 with a stride of 2.

Table IV compares the results for the two denoiser types for both objects, using the same denoiser training policy as in Section VI-C, and experimental configurations of Table VI in the Supplementary Material Section A. The results show little difference between the patch-based denoiser and full frame-based denoiser RED-PSM variants, thus verifying the effectiveness of the patch-based version. The analysis of the computational requirements of the patch-based RED-PSM variant in Section IV-D shows its potential for a highly scalable implementation.

We note that in a divergent beam scenario, the contribution of a patch to a projection will be position-dependent and this would require accurate bookkeeping. Details for doing so can be found in tile-based methods for fan-beam [139] and cone-beam [140] tomography. However, we leave such analysis for future studies.

VII. CONCLUSION

We proposed RED-PSM, the first PSM-based approach to dynamic imaging using a pre-trained and learned (RED-based) spatial prior. The objective in the proposed variational formulation is optimized using a novel and effective bi-convex ADMM algorithm, which enforces the PSM as a hard constraint. Unlike existing PSM-based techniques, RED-PSM is supported by theoretical analysis, with a convergence guarantee to a stationary point of the objective. The results of the numerical experiments show better reconstruction accuracy and considerably faster run times compared to a recent DIP-based algorithm. A patch-based regularizer version of RED-PSM provides almost equivalent performance with a massive reduction of storage requirements, indicating the potential of our framework for dynamic high-resolution 2D or 3D settings.

Possible directions for future work include the application of RED-PSM to different imaging scenarios other than tomography and MRI, and robust denoiser training for RED framework since the deep denoisers encounter varying artifact distributions during optimization. This could also improve the generalizability of the framework to different input types.

ACKNOWLEDGMENT

The authors would like to thank Dr. Brian M. Patterson for providing the compressed material experiment data.

REFERENCES

- [1] E. Maire et al., "20 Hz X-ray tomography during an in situ tensile test," *Int. J. Fracture*, vol. 200, pp. 3–12, 2016.
- [2] I. Danad et al., "Static and dynamic assessment of myocardial perfusion by computed tomography," *Eur. Heart J. - Cardiovasc. Imag.*, vol. 17, no. 8, pp. 836–844, 2016.
- [3] P. J. Keall et al., "Acquiring 4D thoracic CT scans using a multislice helical method," *Phys. Med. Biol.*, vol. 49, no. 10, 2004, Art. no. 2053.
- [4] A. Scanziani et al., "In situ characterization of immiscible three-phase flow at the pore scale for a water-wet carbonate rock," *Adv. Water Resour.*, vol. 121, pp. 446–455, 2018.
- [5] J. O'Connor et al., "Dynamic contrast-enhanced imaging techniques: CT and MRI," *Brit. J. Radiol.*, vol. 84, pp. S112–S120, 2011.
- [6] B. M. Patterson et al., "In situ x-ray synchrotron tomographic imaging during the compression of hyper-elastic polymeric materials," *J. Mater. Sci.*, vol. 51, no. 1, pp. 171–187, 2016.
- [7] B. M. Patterson et al., "Synchrotron CT imaging of lattice structures with engineered defects," *J. Mater. Sci.*, vol. 55, no. 25, pp. 11353–11366, 2020.
- [8] M. Salerno et al., "Recent advances in cardiovascular magnetic resonance," *Circulation: Cardiovasc. Imag.*, vol. 10, no. 6, p. e003951, 2017, doi: [10.1161/CIRCIMAGING.116.003951](https://doi.org/10.1161/CIRCIMAGING.116.003951).
- [9] Z.-P. Liang, "Spatiotemporal imaging with partially separable functions," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2007, pp. 988–991.
- [10] J. P. Haldar and Z.-P. Liang, "Spatiotemporal imaging with partially separable functions: A matrix recovery approach," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2010, pp. 716–719.
- [11] B. Zhao et al., "Low rank matrix recovery for real-time cardiac MRI," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2010, pp. 996–999.
- [12] S. Goud, Y. Hu, and M. Jacob, "Real-time cardiac MRI using low-rank and sparsity penalties," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2010, pp. 988–991.
- [13] Y. Romano et al., "The little engine that could: Regularization by denoising (RED)," *SIAM J. Imag. Sci.*, vol. 10, no. 4, pp. 1804–1844, 2017.
- [14] J. Yoo et al., "Time-dependent deep image prior for dynamic MRI," *IEEE Trans. Med. Imag.*, vol. 40, no. 12, pp. 3337–3348, Dec. 2021.
- [15] S. V. Venkatakrisnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Glob. Conf. Signal Inf. Process.*, 2013, pp. 945–948.
- [16] U. S. Kamilov et al., "Plug-and-play methods for integrating physical and learned models in computational imaging: Theory, algorithms, and applications," *IEEE Signal Process. Mag.*, vol. 40, no. 1, pp. 85–97, Jan. 2023.
- [17] X. Xu et al., "Provable convergence of plug-and-play priors with MMSE denoisers," *IEEE Signal Process. Lett.*, vol. 27, pp. 1280–1284, 2020.
- [18] R. Cohen, M. Elad, and P. Milanfar, "Regularization by denoising via fixed-point projection (RED-PRO)," *SIAM J. Imag. Sci.*, vol. 14, no. 3, pp. 1374–1406, 2021.
- [19] D. Xue et al., "On the convergence of non-convex phase retrieval with denoising priors," *IEEE Trans. Signal Process.*, vol. 70, pp. 4424–4439, 2022.
- [20] E. T. Reehorst and P. Schniter, "Regularization by denoising: Clarifications and new interpretations," *IEEE Trans. Comput. Imag.*, vol. 5, no. 1, pp. 52–67, Mar. 2019.
- [21] X. Hu and T. Parrish, "Reduction of field of view for dynamic imaging," *Magn. Reson. Med.*, vol. 31, no. 6, pp. 691–694, 1994.
- [22] M. E. Brummer et al., "Noquist: Reduced field-of-view imaging by direct fourier inversion," *Magn. Reson. Med.*, vol. 51, no. 2, pp. 331–342, 2004.

- [23] B. Madore, G. H. Glover, and N. J. Pelc, "Unaliasing by Fourier-encoding the overlaps using the temporal dimension (UNFOLD), applied to cardiac imaging and fMRI," *Magn. Reson. Med.*, vol. 42, no. 5, pp. 813–828, 1999.
- [24] M. Lustig et al., "K-t SPARSE: High frame rate dynamic MRI exploiting spatio-temporal sparsity," in *Proc. 13th Annu. Meeting ISMRM, Seattle, 2006*, p. 2420.
- [25] U. Gampfer, P. Boesiger, and S. Kozzerke, "Compressed sensing in dynamic MRI," *Magn. Reson. Med.*, vol. 59, no. 2, pp. 365–373, 2008.
- [26] H. Jung, J. C. Ye, and E. Y. Kim, "Improved k-t BLAST and k-t SENSE using FOCUSS," *Phys. Med. Biol.*, vol. 52, no. 11, 2007, Art. no. 3201.
- [27] H. Jung et al., "K-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI," *Magn. Reson. Med.*, vol. 61, no. 1, pp. 103–116, 2009.
- [28] Q. Zhao, N. Aggarwal, and Y. Bresler, "Dynamic imaging of time-varying objects," in *Proc. ISMRM, 2001*, p. 1776.
- [29] N. Aggarwal, Q. Zhao, and Y. Bresler, "Spatio-temporal modeling and minimum redundancy adaptive acquisition in dynamic MRI," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2002, pp. 737–740.
- [30] N. Aggarwal and Y. Bresler, "Patient-adapted reconstruction and acquisition dynamic imaging method (PARADIGM) for MRI," *Inverse Problems*, vol. 24, no. 4, pp. 0450151–04501529, 2008.
- [31] B. Sharif et al., "Patient-adaptive reconstruction and acquisition in dynamic imaging with sensitivity encoding (PARADISE)," *Magn. Reson. Med.*, vol. 64, no. 2, pp. 501–513, 2010.
- [32] N. P. Willis and Y. Bresler, "Lattice-theoretic analysis of time-sequential sampling of spatio-temporal signals part I," *IEEE Trans. Infor. Theory*, vol. 43, no. 1, pp. 190–207, Jan. 1997.
- [33] N. P. Willis and Y. Bresler, "Lattice-theoretic analysis of time-sequential sampling of spatio-temporal signals part II: Large space-bandwidth-product asymptotics," *IEEE Trans. Infor. Theory*, vol. 43, no. 1, pp. 208–220, Jan. 1997.
- [34] Z.-P. Liang and P. Lauterbur, "An efficient method for dynamic magnetic resonance imaging," *IEEE Trans. Med. Imag.*, vol. 13, no. 4, pp. 677–686, Dec. 1994.
- [35] Z.-P. Liang et al., "Dynamic imaging by model estimation," *Int. J. Imag. Syst. Technol.*, vol. 8, no. 6, pp. 551–557, 1997.
- [36] J. Tsao, P. Boesiger, and K. P. Pruessmann, "K-t BLAST and k-t SENSE: Dynamic MRI with high frame rate exploiting spatiotemporal correlations," *Magn. Reson. Med.*, vol. 50, no. 5, pp. 1031–1042, 2003.
- [37] H. Pedersen et al., "k-t PCA: Temporally constrained k-t BLAST reconstruction using principal component analysis," *Magn. Reson. Med.*, vol. 62, no. 3, pp. 706–716, 2009.
- [38] C. Brinegar et al., "Improving temporal resolution of pulmonary perfusion imaging in rats using the partially separable functions model," *Magn. Reson. Med.*, vol. 64, no. 4, pp. 1162–1170, 2010.
- [39] J. P. Haldar and Z.-P. Liang, "Low-rank approximations for dynamic imaging," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2011, pp. 1052–1055.
- [40] B. Zhao et al., "Further development of image reconstruction from highly undersampled (k, t)-space data with joint partial separability and sparsity constraints," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2011, pp. 1593–1596.
- [41] B. Zhao et al., "Image reconstruction from highly undersampled (k, t)-space data with joint partial separability and sparsity constraints," *IEEE Trans. Med. Imag.*, vol. 31, no. 9, pp. 1809–1820, Sep. 2012.
- [42] S. G. Lingala and M. Jacob, "Blind compressive sensing dynamic MRI," *IEEE Trans. Med. Imag.*, vol. 32, no. 6, pp. 1132–1145, Jun. 2013.
- [43] Y. Djebra, T. Marin, P. K. Han, I. Bloch, G. E. Fakhri, and C. Ma, "Manifold learning via linear tangent space alignment (LTSA) for accelerated dynamic MRI with sparse sampling," *IEEE Trans. Med. Imag.*, vol. 42, no. 1, pp. 158–169, Jan. 2023.
- [44] S. G. Lingala et al., "Accelerated dynamic MRI exploiting sparsity and low-rank structure: Kt SLR," *IEEE Trans. Med. Imag.*, vol. 30, no. 5, pp. 1042–1054, May 2011.
- [45] A. Majumdar and R. K. Ward, "An algorithm for sparse MRI reconstruction by Schatten p-norm minimization," *Magn. Reson. Imag.*, vol. 29, no. 3, pp. 408–417, Apr. 2011.
- [46] A. Majumdar, R. K. Ward, and T. Aboulnasr, "Non-convex algorithm for sparse and low-rank recovery: Application to dynamic MRI reconstruction," *Magn. Reson. Imag.*, vol. 31, no. 3, pp. 448–455, 2013.
- [47] A. Majumdar and R. Ward, "Learning space-time dictionaries for blind compressed sensing dynamic MRI reconstruction," in *Proc. IEEE Int. Conf. Image Process.*, 2015, pp. 4550–4554.
- [48] J. Trzasko and A. Manduca, "Local versus global low-rank promotion in dynamic MRI series reconstruction," in *Proc. Int. Soc. Mag. Reson. Med.*, vol. 24, 2011, p. 4371.
- [49] J. D. Trzasko, "Exploiting local low-rank structure in higher-dimensional MRI applications," *Wavelets Sparsity XV*, vol. 8858, pp. 551–558, Sep. 2013.
- [50] S. G. Lingala, E. DiBella, and M. Jacob, "Deformation corrected compressed sensing (DC-CS): A novel framework for accelerated dynamic MRI," *IEEE Trans. Med. Imag.*, vol. 34, no. 1, pp. 72–85, Jan. 2015.
- [51] X. Chen et al., "Motion-compensated compressed sensing for dynamic contrast-enhanced MRI using regional spatiotemporal sparsity and region tracking: Block low-rank sparsity with motion-guidance (BLOSM)," *Magn. Reson. Med.*, vol. 72, no. 4, pp. 1028–1038, 2014.
- [52] H. Yoon et al., "Motion adaptive patch-based low-rank approach for compressed sensing cardiac cine MRI," *IEEE Trans. Med. Imag.*, vol. 33, no. 11, pp. 2069–2085, Nov. 2014.
- [53] A. Tolouee, J. Alirezaie, and P. Babyn, "Nonrigid motion compensation in compressed sensing reconstruction of cardiac cine MRI," *Magn. Reson. Imag.*, vol. 46, pp. 114–120, 2018.
- [54] B. Trémouh ac et al., "Dynamic MR image reconstruction–separation from undersampled (k,t)-space via low-rank plus sparse prior," *IEEE Trans. Med. Imag.*, vol. 33, no. 8, pp. 1689–1701, Aug. 2014.
- [55] M. Chiew et al., "k-t FASTER: Acceleration of functional MRI data acquisition using low rank constraints," *Magn. Reson. Med.*, vol. 74, no. 2, pp. 353–364, 2015.
- [56] R. Otazo, E. Candes, and D. K. Sodickson, "Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components," *Magn. Reson. Med.*, vol. 73, no. 3, pp. 1125–1136, 2015.
- [57] S. Ravishanker et al., "Low-rank and adaptive sparse signal (LASSI) models for highly accelerated dynamic imaging," *IEEE Trans. Med. Imag.*, vol. 36, no. 5, pp. 1116–1128, May 2017.
- [58] S. Babu, S. G. Lingala, and N. Vaswani, "Fast low rank column-wise compressive sensing for accelerated dynamic MRI," *IEEE Trans. Comput. Imag.*, vol. 9, pp. 409–424, 2023.
- [59] J. Schlemper et al., "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 2, pp. 491–503, Feb. 2018.
- [60] C. Qin et al., "Convolutional recurrent neural networks for dynamic MR image reconstruction," *IEEE Trans. Med. Imag.*, vol. 38, no. 1, pp. 280–290, Jan. 2019.
- [61] S. Biswas, H. K. Aggarwal, and M. Jacob, "Dynamic MRI using model-based deep learning and STORM priors: MoDL-STORM," *Magn. Reson. Med.*, vol. 82, no. 1, pp. 485–494, 2019.
- [62] A. Hauptmann et al., "Real-time cardiovascular MR with spatio-temporal artifact suppression using deep learning—proof of concept in congenital heart disease," *Magn. Reson. Med.*, vol. 81, no. 2, pp. 1143–1156, Feb. 2019.
- [63] T. K ustner et al., "CINeNet: Deep learning-based 3D cardiac CINE MRI reconstruction with multi-coil complex-valued 4D spatio-temporal convolutions," *Sci. Rep.*, vol. 10, no. 1, Aug. 2020, Art. no. 13710.
- [64] Q. Huang et al., "Dynamic MRI reconstruction with end-to-end motion-guided network," *Med. Image Anal.*, vol. 68, Feb. 2021, Art. no. 101901.
- [65] Z. Ke et al., "Deep manifold learning for dynamic MR imaging," *IEEE Trans. Comput. Imag.*, vol. 7, pp. 1314–1327, 2021.
- [66] Z. Ke et al., "Learned low-rank priors in dynamic MR imaging," *IEEE Trans. Med. Imag.*, vol. 40, no. 12, pp. 3698–3710, Dec. 2021.
- [67] C. M. Sandino et al., "Accelerating cardiac cine MRI using a deep learning-based ESPIRiT reconstruction," *Magn. Reson. Med.*, vol. 85, no. 1, pp. 152–167, 2021.
- [68] V. Ghodrati et al., "Temporally aware volumetric generative adversarial network-based MR image reconstruction with simultaneous respiratory motion compensation: Initial feasibility in 3D dynamic cine cardiac MRI," *Magn. Reson. Med.*, vol. 86, no. 5, pp. 2666–2683, 2021.
- [69] T. K ustner et al., "LAPNet: Non-rigid registration derived in k-space for magnetic resonance imaging," *IEEE Trans. Med. Imag.*, vol. 40, no. 12, pp. 3686–3697, Dec. 2021.
- [70] H. Qi et al., "End-to-end deep learning nonrigid motion-corrected reconstruction for highly accelerated free-breathing coronary MRA," *Magn. Reson. Med.*, vol. 86, no. 4, pp. 1983–1996, 2021.
- [71] K. Hammernik et al., "Motion-guided physics-based learning for cardiac MRI reconstruction," in *Proc. IEEE Asilomar Conf. Signals, Syst., Comput.*, 2021, pp. 900–907.
- [72] S. Wang et al., "DIMENSION: Dynamic MR imaging with both k-space and spatial prior knowledge obtained via multi-supervised network training," *NMR Biomed.*, vol. 35, no. 4, 2022, Art. no. e4131.
- [73] Q. Zou et al., "Dynamic imaging using a deep generative storm (Gen-STORM) model," *IEEE Trans. Med. Imag.*, vol. 40, no. 11, pp. 3102–3112, Nov. 2021.

- [74] A. H. Ahmed et al., "Dynamic imaging using deep bi-linear unsupervised representation (DEBLUR)," *IEEE Trans. Med. Imag.*, vol. 41, no. 10, pp. 2693–2703, Oct. 2022.
- [75] N. P. Willis and Y. Bresler, "Tomographic imaging of time-varying distributions," *Biomed. Image Process.*, vol. 1245, pp. 111–123, 1990.
- [76] P. Willis and Y. Bresler, "Optimal scan for time-varying tomography. I. Theoretical analysis and fundamental limitations," *IEEE Trans. Image Process.*, vol. 4, no. 5, pp. 642–653, May 1995.
- [77] N. Willis and Y. Bresler, "Optimal scan for time-varying tomography. II. Efficient design and experimental validation," *IEEE Trans. Image Process.*, vol. 4, no. 5, pp. 654–666, May 1995.
- [78] C. Jaiilin and S. Roux, "Dynamic tomographic reconstruction of deforming volumes," *Materials*, vol. 11, no. 8, 2018, Art. no. 1395.
- [79] G. Zang et al., "Space-time tomography for continuously deforming objects," *ACM Trans. Graph.*, vol. 37, no. 4, pp. 1–14, Jul. 2018.
- [80] C. Jaiilin et al., "Projection-based dynamic tomography," *Phys. Med. Biol.*, vol. 66, no. 21, 2021, Art. no. 215018.
- [81] S. Capostagno et al., "Deformable motion compensation for interventional cone-beam CT," *Phys. Med. Biol.*, vol. 66, no. 5, Feb. 2021, Art. no. 055010.
- [82] M. Burger et al., "A variational reconstruction method for undersampled dynamic X-ray tomography based on physical motion models," *Inverse Problems*, vol. 33, no. 12, 2017, Art. no. 124008.
- [83] B. N. Hahn et al., "Using the Navier-Cauchy equation for motion estimation in dynamic imaging," *Inverse Problems Imag.*, vol. 16, no. 5, pp. 1179–1198, 2022.
- [84] L. Brombal et al., "Motion artifacts assessment and correction using optical tracking in synchrotron radiation breast CT," *Med. Phys.*, vol. 48, no. 9, pp. 5343–5355, 2021.
- [85] A. K. George et al., "Time-resolved cardiac CT reconstruction using the ensemble Kalman filter," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2008, pp. 1489–1492.
- [86] M. D. Butala et al., "Tomographic imaging of dynamic objects with the ensemble Kalman filter," *IEEE Trans. Image Process.*, vol. 18, no. 7, pp. 1573–1587, Jul. 2009.
- [87] J. Hakkarainen et al., "Undersampled dynamic X-ray tomography with dimension reduction Kalman filter," *IEEE Trans. Comput. Imag.*, vol. 5, no. 3, pp. 492–501, Sep. 2019.
- [88] X. Cao et al., "Undersampled dynamic tomography with separated spatial and temporal regularization," in *Proc. 7th Int. Conf. Image Formation X-Ray Computed Tomogr.*, 2022, vol. 12304, pp. 288–295.
- [89] K. A. Mohan et al., "TIMBIR: A method for time-space reconstruction from interlaced views," *IEEE Trans. Comput. Imag.*, vol. 1, no. 2, pp. 96–111, Jun. 2015.
- [90] F. Lucka et al., "Enhancing compressed sensing 4D photoacoustic tomography by simultaneous motion estimation," *SIAM J. Imag. Sci.*, vol. 11, no. 4, pp. 2224–2253, Jan. 2018.
- [91] T. A. Bubba et al., "Sparse dynamic tomography: A shearlet-based approach for iodine perfusion in plant stems," *Inverse Problems*, vol. 36, no. 9, Sep. 2020, Art. no. 094002.
- [92] H. Gao et al., "Robust principal component analysis-based four-dimensional computed tomography," *Phys. Med. Biol.*, vol. 56, no. 11, 2011, Art. no. 3181.
- [93] B. Iskender et al., "Dynamic tomography reconstruction by projection-domain separable modeling," in *Proc. IEEE 14th Image, Video, Multidimensional Signal Process. Workshop*, 2022, pp. 1–5.
- [94] B. Iskender et al., "Dynamic tomography reconstruction by projection-domain separable modeling," 2022, *arXiv:2204.09935*.
- [95] B. Iskender et al., "Factorized projection-domain spatio-temporal regularization for dynamic tomography," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2023, pp. 1–5.
- [96] S. Zhi et al., "Artifacts reduction method for phase-resolved cone-beam CT (CBCT) images via a prior-guided CNN," *Med. Imag. Phys. Med. Imag.*, vol. 10948, pp. 562–567, Mar. 2019.
- [97] X. Huang et al., "U-net-based deformation vector field estimation for motion-compensated 4D-CBCT reconstruction," *Med. Phys.*, vol. 47, no. 7, pp. 3000–3012, 2020.
- [98] F. Madesta et al., "Self-contained deep learning-based boosting of 4D cone-beam CT reconstruction," *Med. Phys.*, vol. 47, no. 11, pp. 5619–5631, 2020.
- [99] Z. Zhang et al., "Deep learning-based motion compensation for four-dimensional cone-beam computed tomography (4D-CBCT) reconstruction," *Med. Phys.*, vol. 50, no. 2, pp. 808–820, 2023.
- [100] S. Majee et al., "Multi-slice fusion for sparse-view and limited-angle 4D CT reconstruction," *IEEE Trans. Comput. Imag.*, vol. 7, pp. 448–462, 2021.
- [101] G. T. Buzzard et al., "Plug-and-play unplugged: Optimization-free reconstruction using consensus equilibrium," *SIAM J. Imag. Sci.*, vol. 11, no. 3, pp. 2001–2020, 2018.
- [102] L. Lozenski, M. A. Anastasio, and U. Villa, "A memory-efficient self-supervised dynamic image reconstruction method using neural fields," *IEEE Trans. Comput. Imag.*, vol. 8, pp. 879–892, 2022.
- [103] X. Shi et al., "Convolutional LSTM network: A machine learning approach for precipitation nowcasting," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, vol. 28, pp. 1–9.
- [104] L. Wang et al., "Spatial-temporal transformer for video snapshot compressive imaging," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 7, pp. 9072–9089, Jul. 2023.
- [105] N. Ginio et al., "Efficient machine learning method for spatio-temporal water surface waves reconstruction from polarimetric images," *Meas. Sci. Technol.*, vol. 34, no. 5, 2023, Art. no. 055801.
- [106] B. Iskender, M. L. Klasky, and Y. Bresler, "RED-PSM: Regularization by denoising of partially separable models for dynamic imaging," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 10561–10570.
- [107] S. Ma et al., "Dynamic MRI reconstruction exploiting partial separability and t-SVD," in *Proc. 7th Int. Conf. Bioinf. Comput. Bio.*, 2019, pp. 179–184.
- [108] S. Ma et al., "Dynamic MRI exploiting partial separability and shift invariant discrete wavelet transform," in *Proc. 6th Int. Conf. Image Vis. Comput.*, 2021, pp. 242–246.
- [109] Y. Sun et al., "Block coordinate regularization by denoising," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, vol. 32, pp. 1–11.
- [110] C. Metzler et al., "prDeep: Robust phase retrieval with a flexible deep network," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 3501–3510.
- [111] G. Mataev, P. Milanfar, and M. Elad, "DeepRED: Deep image prior powered by RED," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshops*, 2019, pp. 1–10.
- [112] S. Boyd et al., "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
- [113] D. Gabay and B. Mercier, "A dual algorithm for the solution of nonlinear variational problems via finite element approximation," *Comput. Math. Appl.*, vol. 2, no. 1, pp. 17–40, 1976.
- [114] J. Eckstein and D. P. Bertsekas, "On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Math. Program.*, vol. 55, no. 1, pp. 293–318, 1992.
- [115] K. Zhang et al., "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017.
- [116] D. Hajinezhad and Q. Shi, "Alternating direction method of multipliers for a class of nonconvex bilinear optimization: Convergence analysis and applications," *J. Glob. Optim.*, vol. 70, pp. 261–288, 2018.
- [117] Y. Wang et al., "Global convergence of ADMM in nonconvex nonsmooth optimization," *J. Sci. Comput.*, vol. 78, pp. 29–63, 2019.
- [118] R. Cohen et al., "It has potential: Gradient-driven denoisers for convergent solutions to inverse problems," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, vol. 34, pp. 18152–18164.
- [119] S. Hurault, A. Leclaire, and N. Papadakis, "Gradient step denoiser for convergent plug-and-play," in *Proc. Int. Conf. Learn. Representations*, 2022, pp. 1–30. [Online]. Available: <https://openreview.net/forum?id=fPhKeld3Okz>
- [120] R. Fermanian, M. L. Pendu, and C. Guillemot, "PnP-ReG: Learned regularizing gradient for plug-and-play gradient descent," *SIAM J. Imag. Sci.*, vol. 16, no. 2, pp. 585–613, 2023.
- [121] A. Berk et al., "Deep proximal gradient method for learned convex regularizers," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2023, pp. 1–5.
- [122] J. R. Chand and M. Jacob, "Multi-scale energy (MuSE) plug and play framework for inverse problems," 2023, *arXiv:2305.04775*.
- [123] M. Udell et al., "Generalized low rank models," *Foundations Trends Mach. Learn.*, vol. 9, no. 1, pp. 1–118, 2016.
- [124] J. D. Rennie and N. Srebro, "Fast maximum margin matrix factorization for collaborative prediction," in *Proc. 22nd Int. Conf. Mach. Learn.*, 2005, pp. 713–719.
- [125] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, 2010.

- [126] H. D. Sarkissian et al., “A cone-beam X-ray computed tomography data collection designed for machine learning,” *Sci. Data*, vol. 6, no. 1, 2019, Art. no. 215.
- [127] S. van der Walt et al., “scikit-image: Image processing in Python,” *PeerJ*, vol. 2, p. e453, Jun. 2014, doi: [10.7717/peerj.453](https://doi.org/10.7717/peerj.453)
- [128] B. M. Patterson, “Personal communication,” 2022.
- [129] F. Reda et al., “FILM: Frame interpolation for large motion,” in *Proc. Eur. Conf. Comput. Vis.*, 2022.
- [130] B. Iskender, M. L. Klasky, and Y. Bresler, “Dynamic tomography reconstruction via low-rank modeling with a red spatial prior,” in *Proc. IEEE 9th Int. Workshop Comput. Adv. Multi-Sensor Adaptive Process.*, 2023, pp. 506–510.
- [131] J. Yoo, TDDIP, GitHub repository, 2021. [Online]. Available: <https://github.com/jaejun-yoo/TDDIP>
- [132] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. Int. Conf. Learn. Representations*, San Diego, CA, USA, 2015.
- [133] O. Bernard et al., “Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: Is the problem solved?,” *IEEE Trans. Med. Imag.*, vol. 37, no. 11, pp. 2514–2525, Nov. 2018.
- [134] S. Kulkarni, J. A. Rumberger, and S. Jha, “Electron beam CT: A historical review,” *Amer. J. Roentgenol.*, vol. 216, no. 5, pp. 1222–1228, 2021.
- [135] D. Spronk et al., “Evaluation of carbon nanotube X-ray source array for stationary head computed tomography,” *Med. Phys.*, vol. 48, no. 3, pp. 1089–1099, 2021.
- [136] S. Xu et al., “Volumetric computed tomography with carbon nanotube X-ray source array for improved image quality and accuracy,” *Commun. Eng.*, vol. 2, no. 1, pp. 1–9, Oct. 2023.
- [137] Z. Wang et al., “Image quality assessment: From error visibility to structural similarity,” *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [138] S. Ravishanker and Y. Bresler, “MR image reconstruction from highly undersampled K-space data by dictionary learning,” *IEEE Trans. Med. Imag.*, vol. 30, no. 5, pp. 1028–1041, May 2011.
- [139] S. Xiao, Y. Bresler, and D. C. Munson, “ $O(n/\log n)$ native fan-beam tomographic reconstruction,” in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2002, pp. 824–827.
- [140] Y. Bresler and J. Brokish, “Fast hierarchical backprojection for helical cone-beam tomography,” in *Proc. IEEE Int. Conf. Image Process.*, 2003, vol. 2, pp. II–815.
- [141] S. Zlobec, “On the Liu–Floudas convexification of smooth programs,” *J. Glob. Optim.*, vol. 32, no. 3, pp. 401–407, 2005.



Berk Iskender (Graduate Student Member, IEEE) received the B.Sc. degree in electrical and electronics engineering from Middle East Technical University, Ankara, Türkiye, in 2018, and the M.S. degree in electrical and computer Engineering in 2020 from the University of Illinois at Urbana-Champaign, Champaign, IL, USA, where he is currently working toward the Ph.D. degree. His research interests include computational imaging, signal processing, computer vision, and machine learning.



Marc L. Klasky is currently a Senior Scientist with the Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM, USA. He has been working in the development and analysis of tomographic methods for several decades in support of the hydrodynamic program. His current research interests include the development of physics-based machine learning algorithms to enable parameter estimation in 3D hydrodynamic simulation codes and the development of reduced-order models to describe the evolution of hydrodynamic features.



Yoram Bresler (Life Fellow, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from the Technion–Israel Institute of Technology, Haifa, Israel, and the Ph.D. from Stanford University, Stanford, CA, USA. He is currently a Professor with the Department of Bioengineering and with the Coordinated Science Laboratory, and Founder Professor Emeritus, Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Champaign, IL, USA. He is also Founder of InstaRecon, Inc., commercializing breakthrough technology for tomographic reconstruction developed in his academic research. His research interests include machine learning and statistical signal processing and their applications to inverse problems in imaging, including compressed sensing, computed tomography, magnetic resonance imaging, and ultrasound. Dr. Bresler is a fellow of the IAMBE, and the AIMBE. His papers have received four IEEE best journal paper awards, two of these with his students. During 2016–2017 he was an IEEE SPS Distinguished Lecturer. He was the recipient of the 1991 NSF Presidential Young Investigator Award, the Technion Fellowship in 1995, and the Xerox Senior Award for Faculty Research in 1998. He was named University of Illinois Scholar in 1999, and was appointed as an Associate with the Center for Advanced Study of the University during 2001–2002. He was a Faculty Fellow with the National Center for Super Computing Applications in 2006. He has served on the editorial board of several journals, including the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the IEEE JOURNAL OF SELECTED TOPICS IN SIGNAL PROCESSING, *Machine Vision and Applications*, and the *SIAM Journal on Imaging Science*, and on various committees of the IEEE.