

Learning Spatially Adaptive ℓ_1 -Norms Weights for Convolutional Synthesis Regularization

Andreas Kofler

Physikalisch-Technische Bundesanstalt (PTB)
Braunschweig and Berlin, Germany
andreas.kofler@ptb.de

Christoph Kolbitsch

Physikalisch-Technische Bundesanstalt (PTB)
Braunschweig and Berlin, Germany
christoph.kolbitsch@ptb.de

Luca Calatroni

MaLGA Center, DIBRIS,
Università di Genova,
MMS, Istituto Italiano di Tecnologia,
Genoa, Italy
luca.calatroni@unige.it

Kostas Papafitsoros

School of Mathematical Sciences
Queen Mary University of London
London, UK
k.papafitsoros@qmul.ac.uk

Abstract—We propose an unrolled algorithm approach for learning spatially adaptive parameter maps in the framework of convolutional synthesis-based ℓ_1 regularization. More precisely, we consider a family of pre-trained convolutional filters and estimate deeply parametrized spatially varying parameters applied to the sparse feature maps by means of unrolling a FISTA algorithm to solve the underlying sparse estimation problem. The proposed approach is evaluated for image reconstruction of low-field MRI and compared to spatially adaptive and non-adaptive analysis-type procedures relying on Total Variation regularization and to a well-established model-based deep learning approach. We show that the proposed approach produces visually and quantitatively comparable results with the latter approaches and at the same time remains highly interpretable. In particular, the inferred parameter maps quantify the local contribution of each filter in the reconstruction, which provides valuable insight into the algorithm mechanism and could potentially be used to discard unsuited filters.

Index Terms—Neural Networks, Convolutional Dictionary Learning, Sparsity, Adaptive Regularization, Low-Field MRI

I. INTRODUCTION

Convolutional Sparse Coding (CSC) and Convolutional Dictionary Learning (CDL) approaches [1] rely on the assumption that the signal/image of interest $\mathbf{x} \in \mathbb{R}^N$ (or $\mathbb{C}^N$) is well-approximated by a linear combination of convolutions of normalized filters $d_k \in \mathbb{R}^{k_f \times k_f}$ with sparse feature maps $\mathbf{s}_k \in \mathbb{R}^N$ (or $\mathbb{C}^N$) for $k = 1, \dots, K$, that is

$$\mathbf{x} \approx \sum_{k=1}^K d_k * \mathbf{s}_k, \quad \text{where } \forall k \text{ } \mathbf{s}_k \text{ is sparse.} \quad (1)$$

The project (22HLT02 A4IM) has received funding from the European Partnership on Metrology, co-financed from the European Union's Horizon Europe Research and Innovation Programme and by the Participating States. LC acknowledges the financial support of the European Research Council (grant MALIN, 101117133).

By enforcing sparsity by means of ℓ_1 -regularization, given a set of signals $\{\mathbf{x}_l\}_{l=1}^L$ and some $\lambda > 0$, the CDL problem can be thus typically formulated as:

$$\min_{\{d_k\}, \{\mathbf{s}_k\}} \frac{1}{2} \sum_{l=1}^L \|\mathbf{x}_l - \sum_{k=1}^K d_k * \mathbf{s}_{k,l}\|_2^2 + \lambda \sum_{l=1}^L \sum_{k=1}^K \|\mathbf{s}_{k,l}\|_1 \quad (2)$$

such that $\|d_k\|_2 = 1, \quad \forall k = 1, \dots, K.$

In the context of image reconstruction, given possibly incomplete observed data $\mathbf{y} \in V$ modeled as the noisy output of a linear forward operator $\mathbf{A} : V \rightarrow W$, CDL and CSC are employed as variational regularization methods, see e.g. [2]:

$$\min_{\mathbf{x}, \{d_k\}, \{\mathbf{s}_k\}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \frac{\alpha}{2} \sum_{k=1}^K \|\mathbf{x} - d_k * \mathbf{s}_k\|_2^2 + \lambda \sum_{k=1}^K \|\mathbf{s}_k\|_1 \quad (3)$$

such that $\|d_k\|_2 = 1, \quad \forall k = 1, \dots, K,$

where $\alpha > 0$ enforces the synthesis constraint. Note that in (3) it is implicitly assumed that the underlying sparsifying model is unknown so that the elements $\mathbf{x}, \{d_k\}, \{\mathbf{s}_k\}$, $k = 1, \dots, K$ are reconstructed jointly, typically by alternating the minimization between solving (2) and computing the current update of desired image depending on the estimated dictionary. These methods are thus often referred to as *blind* Compressed Sensing approaches, see, e.g. [3] for MRI applications.

Problem (3) is a non-convex optimization problem whose solution is, typically, computationally demanding, and further requires careful tuning of the regularization parameters λ , α , initializations $\mathbf{x}^0, \{d_k^0\}, \{\mathbf{s}_k^0\}$ and, depending on the algorithms considered, possibly also of other parameters coupling the image update and the CSC and CDL, see, e.g. [4] for ADMM schemes. Furthermore, the regularization parameter λ dictates the strength of the imposed regularization in terms of the sparsity of the feature maps only globally and

independently on the specific filter considered. This can be limiting, as the filter-dependent representation of local image details and structures may vary, asking for adaptive weighting of the feature maps, similarly as in [5] where this issue is addressed within a non-convolutional setting (i.e. patch-based dictionary learning and sparse coding).

Contribution. In this work, we propose a synthesis-based spatially adaptive regularization scheme. We structure our proposal around two main modules. First, we learn adaptive parameter maps that weight the convolutional sparse approximation using pre-learned filters. Then, we solve the corresponding weighted- ℓ_1 problem using the fast iterative shrinkage algorithm (FISTA) [6]. By utilizing pre-computed filters, our formulation avoids the need for alternating minimization while remaining sample-adaptive, as the regularization strength is estimated at inference for each specific sample. Our work is inspired by a recently proposed approach [7] which uses adaptive spatio-temporal Total Variation (TV) regularization maps parametrized by convolutional neural networks (CNNs), which are learned in a supervised fashion through unrolling of the primal-dual hybrid-gradient (PDHG) algorithm [8], [9]; see also [10] for a recent extension to Total Generalized Variation. Note that employing standard finite-difference discretization filters of the gradient operator corresponds to enforcing a sparsity-based approach with fixed hand-crafted filters within an analysis setting.

We report several numerical results for exemplar low-field MRI data, demonstrating quantitative improvements over non-adaptive models and providing quantitative indications (encoded in the estimated maps) about the adequacy of individual filters for signal representation.

II. METHODS

We provide details on the chosen problem formulation, the reconstruction algorithm used and the neural networks used to estimate the parameter maps.

A. Problem Formulation

We assume to have access to a set of pre-trained unit-norm convolutional dictionary filters $\{d_k\}_{k=1}^K$ with $d_k \in \mathbb{R}^{k_f \times k_f}$ for all k . Instead of relying on the use of a quadratic term weighted by α to enforce (1) as in (3), we explicitly enforce (1) in terms of a re-parametrization of the underlying reconstruction problem. To make the notation more compact, we thus define:

$$\mathbf{s} := [\mathbf{s}_1, \dots, \mathbf{s}_K]^T, \quad \mathbf{D}\mathbf{s} := \sum_{k=1}^K d_k * \mathbf{s}_k, \quad \mathbf{B} := \mathbf{A}\mathbf{D}, \quad (4)$$

and then consider the problem

$$\min_{\mathbf{s} \in \mathbb{R}^{NK}} \frac{1}{2} \|\mathbf{B}\mathbf{s} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{s}\|_1, \quad (5)$$

where λ globally dictates the sparsity level of *all* K feature maps $\{\mathbf{s}_k\}_k$. To adapt the regularization strength both locally as well as differently for each $\mathbf{s}_k$, we now consider the adaptive version given by:

$$\min_{\mathbf{s}} \frac{1}{2} \|\mathbf{B}\mathbf{s} - \mathbf{y}\|_2^2 + \|\mathbf{A}\mathbf{s}\|_1, \quad (6)$$

with $\mathbf{A} \in \mathbb{R}^{NK \times NK}$ being a block-diagonal operator with strictly positive entries, i.e. $\mathbf{A} := \text{diag}(\mathbf{A}_1, \dots, \mathbf{A}_K)$, so that the term $\|\mathbf{A}\mathbf{s}\|_1$ can be written as $\|\mathbf{A}\mathbf{s}\|_1 := \sum_{k=1}^K \|\mathbf{s}_k\|_{\mathbf{A}_k,1} := \sum_{k=1}^K \sum_{j=1}^N |\mathbf{A}_k[j] \cdot \mathbf{s}_k[j]|_1$. To compute such adaptive ℓ_1 maps, we will make use of an algorithm unrolling strategy. Note that in [11] similar ideas for adaptive least-square regularizers are used.

B. Algorithm Unrolling

To learn the parameter maps $\{\mathbf{A}_k\}_{k=1}^K$, we unroll T iterations of the well-known accelerated proximal gradient descent method FISTA [6]. In particular, we use the version of FISTA presented in [12], which ensures the convergence of the iterates to a minimizer of problem (6). Note that in the FISTA update, the proximal operator of $\tau g(\mathbf{s}) := \tau \|\mathbf{A}\mathbf{s}\|_1$ with $\tau > 0$ denoting the algorithmic step-size, can be explicitly computed by $\text{prox}_g(\mathbf{z}) = \mathcal{S}_{\tau\mathbf{A}}(\mathbf{z})$, the componentwise soft-thresholding operator with thresholds defined by the respective values of the corresponding $\mathbf{A}$ -map. Furthermore, convergence is guaranteed under the choice $\tau = 1/\|\mathbf{B}\|^2$. As an alternative, iterative reweighted- ℓ_1 strategies can be used as in [13], with the disadvantage of being less interpretable from the underlying variational perspective.

C. High-Pass Filtering

Most uses of sparse coding approaches require the application of high-pass filtering for better performance, see e.g. [14]. We thus consider a pre-processing of the data aiming to learn the $\{\mathbf{A}_k\}_k$ -maps for the sparse representation of a high-passed filtered representation of the sought image. For that, we employ simple filtering based on gradient regularization as described in [14]. Namely, for an initial image estimate $\mathbf{x}_0$, we define the β -high-passed image $\mathbf{x}_{\text{high}}$ by $\mathbf{x}_{\text{high}} := \mathbf{x}_0 - \mathbf{x}_{\text{low}}$ where the low-pass component $\mathbf{x}_{\text{low}}$ is obtained by

$$\mathbf{x}_{\text{low}} := \underset{\mathbf{x}}{\text{argmin}} \frac{1}{2} \|\mathbf{x} - \mathbf{x}_0\|_2^2 + \frac{\beta}{2} \|\nabla \mathbf{x}\|_2^2, \quad (7)$$

via solving a linear system, by means, e.g., of the conjugate gradient (CG) method. Then, we enforce (1) on the image $\mathbf{x}_{\text{high}}$, define $\mathbf{y}' := \mathbf{y} - \mathbf{A}\mathbf{x}_{\text{low}}$ and consider

$$\min_{\mathbf{s}} \frac{1}{2} \|\mathbf{B}\mathbf{s} - \mathbf{y}'\|_2^2 + \|\mathbf{A}\mathbf{s}\|_1, \quad (8)$$

where thus the convolutional dictionary $\mathbf{D}$ is required to approximate the high-passed component only, i.e. $\mathbf{D}\mathbf{s} \approx \mathbf{x}_{\text{high}}$ and the term $\mathbf{A}\mathbf{x}_{\text{low}}$ is absorbed in $\mathbf{y}'$.

Further, in contrast to [7], we enforce the parameter maps $\{\mathbf{A}_k\}_k$ to have a component-wise upper bound to avoid the scaling ambiguity between the optimization variables $\mathbf{s}$ and the parameter maps $\{\mathbf{A}_k\}_k$.

D. Network Architecture

To obtain the $\{\mathbf{A}_k\}_k$ -maps, we follow [7]. Our network consists of three blocks: i) a high-pass pre-processing to construct the raw data $\mathbf{y}'$ in (8), ii) a deep-network module which outputs the K parameter maps $\{\mathbf{A}_k\}_k$ from a given




Fig. 1. The proposed pipeline consists of three blocks: the first one is a high-pass pre-processing to construct the raw data $\mathbf{y}'$. The second defines a re-parametrization of adaptive $\{\Lambda_k\}_k$ -maps in terms of a CNN NET_Θ given the input image $\mathbf{x}_0 := \mathbf{A}^H \mathbf{y}$. The third consists of unrolling T iterations of the FISTA algorithm for solving (8). Note that the pre-processing block includes a parameter $\beta > 0$, which is learned together with the network weights Θ .

initial image and iii) the unrolling of T iterations of FISTA to approximately solve problem (8).

The parameter maps $\{\Lambda_k\}_k$ are re-parametrized as the output of an expressive CNN with K output-channels which takes as input an initial image estimate, e.g. $\mathbf{x}_0 := \mathbf{A}^H \mathbf{y}$ or $\mathbf{x}_0 := \mathbf{A}^\dagger \mathbf{y}$. In particular, upon vectorization of 2D images into vectors for $k = 1, \dots, K$, we have that:

$$\Lambda_k := (\text{NET}_\Theta(\mathbf{x}_0))_k := t \cdot \phi \circ (\mathbf{u}_\Theta(\mathbf{x}_0))_k, \quad (9)$$

where $\mathbf{u}_\Theta$ denotes a 2D U-Net [15] with parameters Θ and K (vectorized) output channels, ϕ is the sigmoid activation function acting on all output channels component-wise, and $t > 0$ is a hyper-parameter controlling the upper bound of the obtainable $\{\Lambda_k\}_k$ -maps.

For constructing $\mathbf{y}'$ to be used in the FISTA-block as described in II-C, a CG-module solving (7) is applied. The parameter β in (7) can thus be learned as well.

The third block unrolls T iterations of FISTA to estimate the solution $\mathbf{s}^*$ of problem (8), from which then, an estimate of the image can be retrieved as $\mathbf{x}^* := \mathbf{D}\mathbf{s}^* + \mathbf{x}_{\text{low}}$. Figure 1 illustrates the entire pipeline, which we denote as $\mathcal{N}_\Theta^T$.

III. EXPERIMENTS

A. Dataset

We apply the proposed method to 2D low-field brain MR imaging. Low-field MRI suffers from high noise and low image resolution due to hardware constraints. In contrast to standard MRI, only a single receiver coil is used for data acquisition. Using the notation above, we set here $V = W = \mathbb{C}^N$, $\mathbf{A} := \mathbf{S}_I \mathbf{F}$, where $\mathbf{F} : \mathbb{C}^N \rightarrow \mathbb{C}^N$ denotes the 2D Fourier-transform and $\mathbf{S}_I : \mathbb{C}^N \rightarrow \mathbb{C}^N$ denotes a binary mask that only retains k -space coefficients indexed by $I \subset J = \{1, \dots, N\}$ and masks out measurements corresponding to high frequencies leading to low-resolution images. The measurements $\mathbf{y} \in \mathbb{C}^N$ are corrupted by additive white Gaussian noise realizations $\mathbf{e} \in \mathbb{C}^N$ with standard deviation σ . This problem is ill-posed due to the sub-sampling mask $\mathbf{S}_I$ and the low SNR given by the presence of noise

which, in practice, is typically relatively high. For realistic experiments, we thus chose $\sigma \in \{0.075, 0.15, 0.3\}$.

We rely on supervised training to train the proposed network. For training, validation, and testing, we used a total of 6917 brain MR images taken from the fastMRI dataset [16]. We split these into 4875/1393/649 images for training, validation, and testing respectively. To obtain complex-valued images, coil-sensitivity maps were estimated from the multi-coil k -space data, from which a coil-combined image was reconstructed. For this step, we used the method in [17] implemented in MRpro [18].

B. Comparisons

To compare our results, we considered the well-known analysis-type TV-regularized problem with scalar regularization parameter (TV- λ) as well as its weighted extension where spatially varying regularization parameter maps are learned as in [7] (TV- Λ). As a further comparison, we considered the model-based deep learning (MoDL) approach [19], which can be derived by learning the proximal operator of the regularizer by means of a variable splitting method and a deep CNN parametrization. Note that our approach is more interpretable than MoDL, since, once the filters $\{d_k\}_k$ and the $\{\Lambda_k\}_k$ -maps are learned, the reconstruction computed is a minimizer to the functional (8). The use of the variant of FISTA in [12] guarantees that convergence is guaranteed at a numerical level too. Analogously to TV- Λ , our CDL approach is referred to as CDL- Λ , while CDL- λ refers to the approach used to solve problem (5) with one global (non-adaptive) scalar $\lambda > 0$. We evaluate the results in terms of the structural similarity index measure (SSIM) and peak signal-to-noise ratio (PSNR), which are computed after constraining the images on regions excluding the background where no MR signal is present by means of adequately defined binary masks derived from the target images.

C. Training and Implementation Details

The 2D U-Net $\mathbf{u}_\Theta$ used in the CNN-block NET_Θ in (9) has a 2- K -2 K -4 K -2 K - K structure with two input-channels




Fig. 2. An example of 12 out of $K = 64$ $\{\Lambda_k\}_k$ -maps for dictionary filters with filter size $k_f \times k_f = 11 \times 11$ along with the corresponding sparse feature maps. The first two rows show the feature maps of the real and imaginary part of the solution of (8), while the third row shows the corresponding $\{\Lambda_k\}_k$ -maps. Results are ordered in a decreasing order w.r.t. the variance of the $\{\Lambda_k\}_k$ -maps. The results indicate that the network NET_Θ can provide structural information associated with each filter and, moreover, can identify filters that are less relevant for the reconstruction (high values of the map).




Fig. 3. An example of images reconstructed with the different methods for a noise level of $\sigma = 0.3$. The SSIM and the PSNR, which are computed by masking out values outside of the defined mask, where no signal is present, are listed in the point-wise error images. To better highlight residual errors that can be attributed to the model limitations, the point-wise error images were scaled by a factor of 3.

(for the real and imaginary parts of the image) and K output channels for K parameter maps $\{\Lambda_k\}_k$. To handle complex-valued data, note that the real and the imaginary parts of the images are approximated using the same filters, i.e. we have $\mathbf{D}\mathbf{s} := \sum_k (d_k * \text{Re}(\mathbf{s}_k) + i d_k * \text{Im}(\mathbf{s}_k))$. For training, we used the Adam optimizer [20] with an MSE-loss using a batch-size of 2. For all methods, the scalar regularization parameters (λ and β), which were "activated" using a Softplus activation function to maintain strict positivity, were assigned a learning rate of 10^{-1} , while the U-Nets $\mathbf{u}_\Theta$ were given a learning rate of 10^{-4} with a weight decay regularization weight of 10^{-5} . The dictionary filters $\{d_k\}_k$ were pre-trained on 160 randomly selected images from the training set using the algorithm described in [21] implemented in SPORCO [22]. We tested different kernel sizes $k_f \times k_f = 7 \times 7$, 11×11 and 13×13 as well as different numbers of filters $K = 32, 64$. Due to space limitations, we here only report results for $K = 64$ and $k_f \times k_f = 11 \times 11$. We point out that the experiments suggest that more and larger filters tend to improve the achievable performance. The upper bound $t > 0$ in (9) was fixed to $t = 10.0$, which is also the average maximal magnitude value the images were normalized to.

The number of unrolling steps was set to $T = 128$ for TV- λ and TV- Λ , to $T = 64$ for CDL- λ and CDL- Λ (which enjoy faster convergence inherited from FISTA), and to $T = 8$ for MoDL. The number of epochs was 96 for MoDL (≈ 16 h training), 4 for TV- λ and CDL- λ (≈ 4 h training time) 48 for TV- Λ (≈ 11 h training time) and CDL- Λ (≈ 42 h training time). Upon integration of all necessary building blocks in MRpro [18], the code of this work will be made available at github.com/koflera/LearningL1NormsWeights4Synthesis.

IV. RESULTS AND DISCUSSION

In Figure 2 we see an example of 12 sparse feature maps estimated by FISTA as well as the corresponding learned regularization parameter maps and their associated filters. From the parameter maps, we see that the network seems to be able to distinguish between more and less useful filters for the sparse approximation. Note that higher values in the $\{\Lambda_k\}_k$ -maps result in a higher threshold in the proximal operator of g and thus to a lower contribution in the sparse representation.

Figure 3 shows an example of images reconstructed for a noise level of $\sigma = 0.3$. All methods successfully improved upon the adjoint-based reconstruction, which heavily suffers




Fig. 4. A summary of SSIM and PSNR obtained on the test set by the different approaches. The proposed CDL- λ achieves competitive results compared to the TV- λ [7] as well as to the deep learning-based method MoDL [19].

from noise and blurring artifacts. The methods employing only scalar regularization values λ exhibit larger errors at the boundaries and are clearly outperformed by both TV- λ [7] and the proposed CDL- λ , which employ learned regularization parameter maps and yield results comparable to the deep learning-based method MoDL [19]. Although similar in terms of SSIM and PSNR, the investigated CDL- λ yields more anatomically accurate images in contrast to TV- λ , which, despite the spatially adaptive regularization, tends to exhibit some piece-wise constant structures. Finally, CDL- λ exhibits a high visual similarity when compared to MoDL.

Figure 4 summarizes the results obtained on the test set in terms of SSIM and PSNR. The methods involving only scalar regularization parameters (TV- λ and CDL- λ) are outperformed by both TV- λ and CDL- λ involving the learned adaptive spatial regularization, which are competitive even when compared to the deep learning-based MoDL [19].

V. CONCLUSION

In this work, we extended the method proposed in [7] for learning TV adaptive maps via CNN parametrization and algorithmic unrolling to the sparse convolutional synthesis setting. The proposed approach combines model-based regularization, in our case CSC and CDL, with deep neural networks; it yields competitive results compared to MoDL [19], but further enjoys interpretability, convergence guarantees, and informative quantitative estimation of the local behavior with respect to a pre-trained set of convolutional filters. As a byproduct, the learned parameter maps can be used to discriminate between more and less useful convolutional filters in the representation of the solution, resulting in an at-inference-time self-adaptation of the reconstruction method. As a consequence, this approach offers a valuable alternative to the well-established paradigm of learning new filters jointly with the desired solution for each reconstruction problem at hand.

A possible future direction of this work could be the use

of self-supervised training procedures to estimate adequate regularization maps without the need for reference data.

REFERENCES

- [1] C. Garcia-Cardona and B. Wohlberg, "Convolutional dictionary learning: A comparative review and new algorithms," *IEEE Trans. Comput. Imag.*, vol. 4, no. 3, pp. 366–381, 2018.
- [2] T. M. Quan and W.-K. Jeong, "Compressed sensing reconstruction of dynamic contrast enhanced MRI using GPU-accelerated convolutional sparse coding," in *2016 IEEE 13th ISBI*. IEEE, 2016, pp. 518–521.
- [3] S. G. Lingala and M. Jacob, "Blind compressive sensing dynamic MRI," *IEEE Trans. Med. Imaging*, vol. 32, no. 6, pp. 1132–1145, 2013.
- [4] B. Wohlberg, "Efficient algorithms for convolutional sparse representations," *IEEE Trans. Image Process.*, vol. 25, no. 1, pp. 301–315, 2015.
- [5] M.-C. Pali, T. Schaeffter, C. Kolbitsch, and A. Kofler, "Adaptive sparsity level and dictionary size estimation for image reconstruction in accelerated 2D radial cine MRI," *Med. Phys.*, vol. 48, no. 1, 2021.
- [6] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imaging Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [7] A. Kofler, F. Altekruiger, F. Antaro Ba, C. Kolbitsch, E. Papoutsellis, D. Schote, C. Sirotenko, F. F. Zimmermann, and K. Papafitsoros, "Learning regularization parameter-maps for variational image reconstruction using deep neural networks and algorithm unrolling," *SIAM J. Imaging Sci.*, vol. 16, no. 4, pp. 2202–2246, 2023.
- [8] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Process. Mag.*, vol. 38, no. 2, pp. 18–44, 2021.
- [9] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *J. Math. Imaging Vis.*, vol. 40, pp. 120–145, 2011.
- [10] T. T. Vu, A. Kofler, and K. Papafitsoros, "Deep unrolling for learning optimal spatially varying regularisation parameters for total generalised variation," in *Scale Space and Variational Methods in Computer Vision*. Springer Nature Switzerland, 2025, pp. 282–294.
- [11] M. Pourya, E. Kobler, M. Unser, and S. Neumayer, "Dealing with image reconstruction: Deep attentive least squares," 2025. [Online]. Available: <https://arxiv.org/abs/2502.04079>
- [12] A. Chambolle and C. H. Dossal, "On the convergence of the iterates of "FISTA"," *J. Optim. Theory Appl.*, vol. 166, no. 3, p. 25, 2015.
- [13] E. J. Candès, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *J. Fourier Anal. Appl.*, vol. 14, no. 5, pp. 877–905, 2008.
- [14] B. Wohlberg, "Convolutional sparse representations as an image model for impulse noise restoration," in *2016 IEEE 12th IVMSW workshop*. IEEE, 2016, pp. 1–5.
- [15] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *18th International MICCAI 2015 conference, part III 18*. Springer, 2015, pp. 234–241.
- [16] J. Zbontar, F. Knoll, A. Sriram, T. Murrell, Z. Huang, M. J. Muckley, A. Defazio, R. Stern, P. Johnson, M. Bruno *et al.*, "fastMRI: An open dataset and benchmarks for accelerated MRI," *arXiv preprint:1811.08839*, 2018.
- [17] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin, "Adaptive reconstruction of phased array MR imagery," *Magn. Reson. Med.*, vol. 43, no. 5, pp. 682–690, 2000.
- [18] F. F. Zimmermann, P. Schuenke, S. Brahma, M. Guastini, J. Hammacher, A. Kofler, C. Kranich Redshaw, L. Lunin, S. Martin, D. Schote, and C. Kolbitsch, "MRpro - PyTorch-based MR image reconstruction and processing package," Feb. 2025. [Online]. Available: <https://doi.org/10.5281/zenodo.14886714>
- [19] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-based deep learning architecture for inverse problems," *IEEE Trans. Med. Imaging*, vol. 38, no. 2, pp. 394–405, 2018.
- [20] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [21] J. Liu, C. Garcia-Cardona, B. Wohlberg, and W. Yin, "First-and second-order methods for online convolutional dictionary learning," *SIAM J. Imaging Sci.*, vol. 11, no. 2, pp. 1589–1628, 2018.
- [22] B. Wohlberg, "SPORCO: A python package for standard and convolutional sparse representations," in *SciPy*, 2017, pp. 1–8.