

Harnessing Side Information for Highly Accelerated MRI

Arda Atalık^{1,2,3}, Sumit Chopra^{3,4}, and Daniel K. Sodickson^{2,3}

¹ NYU Center for Data Science, NY, USA

² Center for Advanced Imaging Innovation and Research (CAI²R), Department of Radiology, NYU Grossman School of Medicine, NY, USA

³ Bernard and Irene Schwartz Center for Biomedical Imaging, Department of Radiology, NYU Grossman School of Medicine, NY, USA

⁴ Courant Institute of Mathematical Sciences, NY, USA

{Arda.Atalik,Sumit.Chopra,Daniel.Sodickson}@nyulangone.org

Abstract. Reducing MRI scan times can improve patient care and lower healthcare costs. Many acceleration methods are designed to reconstruct diagnostic-quality images from sparse k -space data, via an ill-posed or ill-conditioned linear inverse problem (LIP). To address the resulting ambiguities, it is crucial to incorporate prior knowledge into the optimization problem, e.g., in the form of regularization. Another form of prior knowledge less commonly used in medical imaging is the *side information* obtained from sources other than the current acquisition. In this paper, we present the **Trust-Guided Variational Network (TGVN)**, an end-to-end deep learning framework that effectively and reliably integrates side information into LIPs. We demonstrate its effectiveness in multi-coil, multi-contrast MRI reconstruction, where incomplete or low-SNR measurements from one contrast are used as side information to reconstruct high-quality images of another contrast from heavily undersampled data. TGVN is robust across different contrasts, anatomies, and field strengths. Compared to baselines utilizing side information, TGVN achieves superior image quality while preserving subtle pathological features even at challenging acceleration levels, drastically speeding up acquisition while minimizing hallucinations. Source code and dataset splits are available on github.com/sodicksonlab/TGVN.

Keywords: MRI reconstruction · Side information · Deep learning

1 Introduction

Magnetic Resonance Imaging (MRI) is a mainstay of medical diagnostic imaging, thanks to its flexibility, its rich information content, and its excellent soft-tissue contrast. An MR scanner collects measurements in frequency space (a.k.a. k -space) that encode the body’s response to applied electromagnetic fields, typically with multiple receiver coils capturing distinct views modulated by their

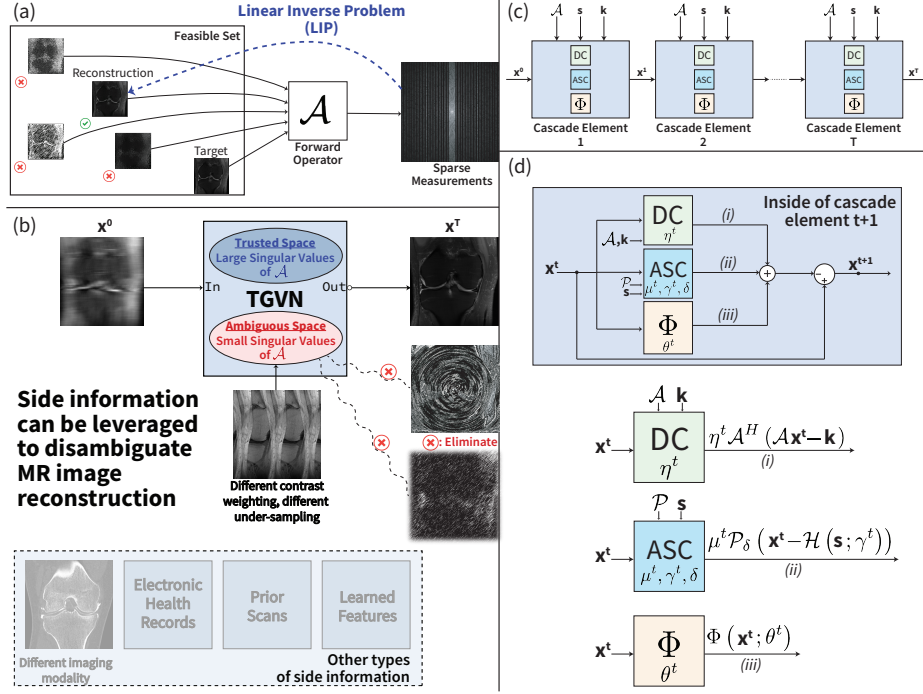


Fig. 1: **TGVN reconstruction with side information:** (a) Visual representation of an MRI reconstruction LIP. (b) Overview of trust-guided disambiguation of solutions to the LIP. We used different contrast-weighted measurements as side information. However, side information can originate from various sources. (c) TGVN: A cascade of T elements connected in series. (d) Components of each element in the cascade—data consistency (DC, parametrized by η), ambiguous space consistency (ASC, parametrized by μ, γ, δ), and refinement (Φ , parametrized by θ)—are shown along with their inputs, outputs, and the final aggregation. Φ and \mathcal{H} are the U-Net [25] models. TGVN’s key novelty is the **ASC module** which incorporates side information \mathbf{s} safely through projection \mathcal{P} onto the ambiguous space of \mathcal{A} without perturbing the trusted space.

individual sensitivities, and this process is mathematically described by a linear map called the *forward operator*. The acquired k -space measurements are then used to reconstruct a spatially resolved image by solving the corresponding linear inverse problem (LIP). Despite MRI’s superior diagnostic capabilities, it is comparatively time-consuming and costly, which limits its overall accessibility. Reducing the time it takes to acquire an MR scan is an important practical problem that can improve patient care by limiting patient discomfort, reducing costs, and improving accessibility of this imaging modality. One way to reduce scan time is to acquire fewer k -space measurements. The challenge then becomes reconstructing high-quality images from limited data by solving the correspond-

ing ill-conditioned/ill-posed LIP, which admits many mathematically feasible solutions, most of which fail to capture essential anatomic and clinical details accurately (see Fig. 1(a)).

Researchers have proposed various solutions, including compressed sensing [20,21] and deep learning (DL) approaches [13,10,1,29,33] which integrate traditional optimization techniques with deep neural networks. Another approach to eliminating degenerate solutions to these LIPs involves leveraging additional contextual information (i.e., *side information*), that is typically readily available in real-world scenarios. The nature of side information is problem-dependent, and it can originate from multiple sources. For example, data associated with prior scans of the same patient can serve as side information, as can images gathered during the same scan but associated with an imaging pulse sequence of different contrast than the target pulse sequence.

In this work, we propose a novel, end-to-end trainable deep learning method that efficiently and reliably integrates side information to solve LIPs, as illustrated in Fig. 1. Our method, called the **Trust-Guided Variational Network (TGVN)**, uses the side information only to disambiguate the subspace of solutions that the forward operator cannot reliably distinguish based on the measured data. We demonstrate the effectiveness of our method in the challenging domain of multi-coil, multi-contrast MRI reconstruction, where incomplete or low-SNR measurements with complementary contrast weighting are used as side information to reconstruct images with a different target contrast from heavily undersampled k -space measurements across different anatomies and field strengths. Compared to recently proposed DL-based solutions, our method leverages side information more efficiently and reliably, preserving fine anatomical and pathological details obliterated by other methods at high acceleration levels, and achieving statistically significant improvements in reconstruction performance.

2 Background

2.1 Deep Learning for Multi-coil MR Image Reconstruction

Let $\tilde{\mathbf{k}}$ denote the fully-sampled k -space measurements, which represent Fourier coefficients of the structure of the continuous object being imaged. We define a discrete estimated MR image \mathbf{x} , such that $\mathbf{k} = \mathcal{F}(\mathbf{x}) + \epsilon$, where ϵ is the measurement noise and \mathcal{F} denotes the Fourier transform operator. In multi-coil acquisition (a.k.a. parallel imaging) [27,23], the scanner captures multiple views of the anatomy modulated by the sensitivities S_i of the receiver coils. Multi-coil k -space measurements are given by $\tilde{\mathbf{k}}_i = \mathcal{F}(S_i \mathbf{x}) + \epsilon_i$, for each $i \in \{1, 2, \dots, N_c\}$, where ϵ_i is the measurement noise for coil i and N_c denotes the number of coils. To simplify notation, we aggregate the k -space data from all coils into a single tensor $\tilde{\mathbf{k}} = (\tilde{\mathbf{k}}_1, \dots, \tilde{\mathbf{k}}_{N_c})$ and define the *expand* operator $\mathcal{E} : \mathbf{x} \mapsto (\mathcal{F}(S_1 \mathbf{x}), \dots, \mathcal{F}(S_{N_c} \mathbf{x}))$. To accelerate MRI acquisition, fewer k -space samples are acquired, which we denote by a binary mask \mathcal{M} . The undersampled k -space measurements can be denoted as $\mathbf{k} \triangleq \mathcal{M}\tilde{\mathbf{k}} = (\mathcal{M}\tilde{\mathbf{k}}_1, \dots, \mathcal{M}\tilde{\mathbf{k}}_{N_c})$, and

the forward operator \mathcal{A} is equal to $\mathcal{M} \circ \mathcal{E}$. That is, $\mathbf{k} = \mathcal{A} \mathbf{x} + \epsilon' = (\mathcal{M} \circ \mathcal{E}) \mathbf{x} + \epsilon'$, where ϵ' denotes the measurement noise in the undersampled k -space.

When \mathbf{k} is undersampled, the LIP of estimating \mathbf{x} from undersampled measurements is often formulated as a regularized least-squares problem, i.e., $\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathcal{A} \mathbf{x} - \mathbf{k}\|_2^2 + \Psi(\mathbf{x})$, where $\Psi(\cdot)$ denote a regularization function that imposes certain constraints on the possible solutions \mathbf{x} .

In DL-based MRI reconstruction methods, such as the End-to-End Variational Network (E2E-VN) [29], one learns a regularization function from the training data to maximize a desired similarity metric between the reconstructed image $\hat{\mathbf{x}}$ and the ground truth. Specifically, E2E-VN starts with an initial estimate \mathbf{x}^0 of the solution to $\mathcal{A} \mathbf{x} = \mathbf{k}$, and uses the Landweber method [16] to iteratively refine its estimate. More formally, starting with $\mathbf{x}^0 = \mathcal{A}^H \mathbf{k}$, it executes the following sequence of steps for a total of T iterations: $\mathbf{x}^{t+1} = \mathbf{x}^t - \eta^t \mathcal{A}^H (\mathcal{A} \mathbf{x}^t - \mathbf{k}) - \Phi(\mathbf{x}^t; \theta^t)$, where \mathcal{A}^H is the Hermitian adjoint of \mathcal{A} , η^t is the step size, and $\Phi(\cdot; \theta^t)$ is a neural network with parameters θ^t , representing the gradient of Ψ . At the end of iteration T , we obtain \mathbf{x}^T parameterized by $\Theta = \{\theta^0, \dots, \theta^{T-1}, \eta^0, \dots, \eta^{T-1}\}$. Assuming access to ground truth \mathbf{x}^* , parameters Θ are learned in a supervised manner to maximize a desired similarity between \mathbf{x}^T and \mathbf{x}^* .

2.2 Related Work

Image reconstruction with side information has a rich history beginning in the 1990s. [8] introduced an emission image reconstruction method that weights a Gibbs prior with MR-derived anatomical boundaries. [9] proposed a Bayesian reconstruction of PET and SPECT with the aid of prior information derived from registered MR images of the same slice. [3] presented a Bayesian method for simultaneously segmenting and reconstructing ECT images utilizing high-resolution anatomical information from other imaging modalities.

Multiple studies have proposed methods that use data from complementary contrasts as side information to accelerate MRI—reconstructing the target contrast by leveraging information from other contrast(s). Some of the earlier attempts utilized handcrafted priors [2,12,6]. Later, dictionary-learning-based methods [28,17] and end-to-end DL-based models were proposed [34,36,7,18]. Recently, generative models utilizing side information were introduced [5,14,19]. Synergizing the range-null space decomposition with DL was first explored in [26], and later works have explored its potential in relevant tasks and domains [30,31,4]. Despite notable advancements, existing methods for solving LIPs that incorporate side information still struggle with highly undersampled data, often leading to degraded image quality or hallucinations. The former can be attributed to a lack of efficiency in exploiting side information (i.e., insufficiency to disambiguate the solution space), while the latter represents over-reliance on it. Consequently, harnessing the full potential of side information while mitigating hallucinations remains an open problem that can have a transformative impact on healthcare.

Table 1: **Experiments:** Each line is a single experiment. Undersampling, in either a random or a uniform pattern, was performed along the phase-encoding (PE) direction. Center frequency (CF) is the portion of fully sampled central PE lines. For knee experiments (K1, K2, and K3), acceleration was achieved solely by undersampling. For brain experiments (B1 and B2), the total acceleration factors include both undersampling ($18\times$ and $15\times$ for main information) and reduction in the number of repetitions ($2\times$ and $3\times$ for main, $3\times$ for side information).

| Exp. | Main information | | | Side information | | |
|------|------------------|---------------------|----|------------------|--------------------|----|
| | Contrast | Acceleration | CF | Contrast | Acceleration | CF |
| K1 | PDFS | Random- $20\times$ | 3% | PD | Uniform- $2\times$ | 0% |
| K2 | PDFS | Random- $14\times$ | 3% | PD | Uniform- $3\times$ | 0% |
| K3 | PDFS | Random- $6\times$ | 5% | PD | Uniform- $3\times$ | 0% |
| B1 | FLAIR | Uniform- $36\times$ | 2% | T2 | $3\times$ | - |
| B2 | T1 | Uniform- $45\times$ | 2% | T2 | $3\times$ | - |

3 Trust Guided Variational Network (TGVN)

Data consistency alone may not suffice to resolve inherent ambiguities in the LIP solution space, particularly at high accelerations where otherwise effective DL reconstruction methods have been observed to fail catastrophically [24]. Here we introduce the concept of *ambiguous space consistency*. We focus on using side information to resolve the ambiguous space characterized by the small singular values of \mathcal{A} , since these components of the problem tend to amplify noise and otherwise introduce uncertainty into prospective solutions.

Let \mathbf{x}_p be a particular solution to $\mathcal{A}\mathbf{x} = \mathbf{k}$ and let $\sum_i \sigma_i \mathbf{u}_i \mathbf{v}_i^H$ represent the singular value decomposition (SVD) of \mathcal{A} . Given a small positive threshold δ , we define the *ambiguous space* as the subspace spanned by the right singular vectors \mathbf{v}_i with corresponding singular values σ_i smaller than δ , and denote it as $\mathcal{W}_\delta(\mathcal{A})$. Observe that if we add any unit vector $\mathbf{x}_a \in \mathcal{W}_\delta(\mathcal{A})$ to \mathbf{x}_p , the data inconsistency loss $\|\mathcal{A}(\mathbf{x}_p + \mathbf{x}_a) - \mathbf{k}\|_2^2$ can at most be δ^2 . In other words, perturbation can create only minor data inconsistency. However, only certain \mathbf{x}_a maximize the desired similarity between $(\mathbf{x}_p + \mathbf{x}_a)$ and \mathbf{x}^* , indicating that, once a particular solution is found, images from $\mathcal{W}_\delta(\mathcal{A})$ introduce ambiguity in the reconstruction problem. Inspired by this observation, we propose to explicitly learn a constraint that removes undesirable solutions from $\mathcal{W}_\delta(\mathcal{A})$.

Our idea is to project \mathbf{x} onto $\mathcal{W}_\delta(\mathcal{A})$ via the orthogonal projector \mathcal{P}_δ to obtain \mathbf{x}_a , and to guide \mathbf{x}_a to be maximally consistent with the side information \mathbf{s} using a learnable module \mathcal{H} parametrized by γ . That is, we add a squared Euclidean distance constraint $\|\mathcal{P}_\delta \mathbf{x} - \mathcal{H}(\mathbf{s}; \gamma)\|_2^2$ to the regularized least-squares formulation to obtain $\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathcal{A}\mathbf{x} - \mathbf{k}\|_2^2 + \frac{\beta}{2} \|\mathcal{P}_\delta \mathbf{x} - \mathcal{H}(\mathbf{s}; \gamma)\|_2^2 + \Psi(\mathbf{x})$. As the added constraint involves only a squared Euclidean distance, its integration into Landweber iterations is straightforward. Starting with $\mathbf{x}_0 = \mathcal{A}^H \mathbf{k}$, we

Table 2: **Quantitative evaluation results:** SSIM [32], PSNR, and NRMSE are reported for knee (K1, K2, K3) and brain (B1, B2) experiments using TGVN and baselines. Each metric shows the mean and standard error over the test set. Bold entries indicate the best performance, consistently achieved by TGVN.

| | Exp. | TGVN | DMSI | MCVN | MTrans | E2E-VN |
|-------|------|----------------------|---------------|----------------------|---------------|---------------|
| SSIM | K1 | 84.92 ± 0.19 | 56.99 ± 0.31 | 82.89 ± 0.21 | 80.84 ± 0.23 | 81.33 ± 0.23 |
| | K2 | 85.52 ± 0.19 | 58.76 ± 0.31 | 83.13 ± 0.21 | 81.25 ± 0.22 | 83.40 ± 0.21 |
| | K3 | 88.02 ± 0.17 | 64.35 ± 0.31 | 86.47 ± 0.18 | 85.29 ± 0.19 | 87.42 ± 0.17 |
| | B1 | 87.34 ± 0.12 | - | 86.95 ± 0.12 | 84.03 ± 0.14 | 81.40 ± 0.15 |
| | B2 | 89.34 ± 0.10 | - | 88.66 ± 0.12 | 85.72 ± 0.17 | 86.11 ± 0.13 |
| PSNR | K1 | 30.92 ± 0.07 | 22.22 ± 0.10 | 29.97 ± 0.07 | 28.93 ± 0.07 | 29.30 ± 0.07 |
| | K2 | 31.31 ± 0.07 | 22.68 ± 0.10 | 30.07 ± 0.07 | 29.11 ± 0.07 | 30.37 ± 0.07 |
| | K3 | 32.89 ± 0.08 | 24.24 ± 0.10 | 31.95 ± 0.07 | 31.26 ± 0.07 | 32.59 ± 0.07 |
| | B1 | 30.81 ± 0.08 | - | 30.75 ± 0.08 | 28.70 ± 0.08 | 27.14 ± 0.08 |
| | B2 | 31.36 ± 0.07 | - | 30.94 ± 0.08 | 28.90 ± 0.10 | 29.31 ± 0.08 |
| NRMSE | K1 | 0.140 ± 0.001 | 0.397 ± 0.004 | 0.157 ± 0.001 | 0.177 ± 0.001 | 0.170 ± 0.001 |
| | K2 | 0.134 ± 0.001 | 0.376 ± 0.004 | 0.155 ± 0.001 | 0.174 ± 0.001 | 0.150 ± 0.001 |
| | K3 | 0.112 ± 0.001 | 0.317 ± 0.003 | 0.124 ± 0.001 | 0.135 ± 0.001 | 0.116 ± 0.001 |
| | B1 | 0.158 ± 0.002 | - | 0.159 ± 0.002 | 0.201 ± 0.002 | 0.240 ± 0.003 |
| | B2 | 0.162 ± 0.002 | - | 0.171 ± 0.002 | 0.218 ± 0.003 | 0.205 ± 0.002 |

execute the following sequence of steps for a total of T iterations.

$$\mathbf{x}^{t+1} = \mathbf{x}^t - \eta^t \mathcal{A}^H (\mathcal{A} \mathbf{x}^t - \mathbf{k}) - \underbrace{\mu^t \mathcal{P}_\delta (\mathbf{x}^t - \mathcal{H}(\mathbf{s}; \gamma^t))}_{\text{trust-guidance}} - \Phi(\mathbf{x}^t; \theta^t). \quad (1)$$

At the end of the iteration T , we obtain \mathbf{x}^T parameterized by $\Omega \triangleq \Theta \cup \{\delta, \gamma^0, \dots, \gamma^{T-1}, \mu^0, \dots, \mu^{T-1}\}$. Assuming access to ground truth \mathbf{x}^* , the parameters Ω are learned in a supervised manner to maximize a desired similarity between \mathbf{x}^T and \mathbf{x}^* . Furthermore, we present an efficient approximation of the exact orthogonal projector \mathcal{P}_δ , which bypasses the need for SVD computation. This approach is crucial for managing the scale of the forward operator, which may contain hundreds of thousands of rows and columns, making explicit methods infeasible. For a set \mathcal{K} , let $1_{\mathcal{K}}(x)$ denote an indicator function that equals 1 if $x \in \mathcal{K}$ and 0 otherwise. Given δ , the exact projector can be written as $\mathcal{P}_\delta = \sum_i 1_{[0, \delta)}(\sigma_i) \mathbf{v}_i \mathbf{v}_i^H$. In lieu of assigning binary weights to the i th projection, we can weight them by $\delta^2 / (\delta^2 + \sigma_i^2)$, and define $\mathcal{P}'_\delta \triangleq \sum_i \frac{\delta^2}{\delta^2 + \sigma_i^2} \mathbf{v}_i \mathbf{v}_i^H = (I + \frac{1}{\delta^2} \mathcal{A}^H \mathcal{A})^{-1}$, where $I : \mathbf{x} \mapsto \mathbf{x}$ denotes the identity operator. We can then calculate the approximate trust-guidance term $\mu^t (I + \frac{1}{\delta^2} \mathcal{A}^H \mathcal{A})^{-1} (\mathbf{x}^t - \mathcal{H}(\mathbf{s}; \gamma^t))$ efficiently using a small number of Conjugate Gradient iterations [11].




Fig. 2: **Using side information significantly enhances reconstruction quality.** **Left:** Coronal PDFS knee reconstructions (K1, $20\times$ acceleration). **Yellow arrows** highlight subtle structures (e.g., meniscal tear, vastus lateralis muscle), preserved only when side information is used. **Right:** Axial FLAIR brain reconstructions (B1, $36\times$ total acceleration: $18\times$ undersampling, $2\times$ repetition reduction). TGVN recovers fine neuroanatomical details (yellow arrows), matching the ground truth, whereas key brain structures are homogenized without it.

4 Empirical Validation and Results

We designed experiments to answer two key questions: Q1) Is there benefit in using side information? Q2) How effective is TGVN at utilizing the side information? To answer Q1, we compared the reconstruction performance of TGVN against E2E-VN of the same capacity that does not utilize side information. Q2 was answered by comparing TGVN’s performance against several DL-based baselines that also leverage side information: MTrans [7], MCVN [18], and DMSI [19]. We also performed ablation studies to assess the impact of the projector \mathcal{P}_δ on reconstruction quality and to document the robustness of TGVN to small random misregistrations between main and side information. While the results of these ablation studies further support the efficacy and reliability of TGVN, we omit these details here due to space constraints.

In knee experiments (K1, K2, K3), we utilized a subset of data from the multi-coil track of the fastMRI knee dataset [15,35], consisting of 368, 30, and 30 volumes for training, validation, and test sets, respectively. Brain experiments (B1, B2) utilized the M4Raw dataset [22]. In all experiments, we used patient-level splits to avoid data leakage. For a fair comparison, TGVN and the baselines were configured to have a similar number of trainable parameters (95M for K1–3, and 67M for B1–2). All models were trained on $4\times$ A100 GPUs for 100 epochs, with parameters saved at each epoch. Each model took up to 10 days to train per knee experiment and 2 days per brain experiment, and the model with the best validation SSIM was selected for testing. Details regarding the experimental settings, including the contrast and acceleration details, are provided in Table 1.

Fig. 2 (left) shows the reconstruction results for coronal PDFS knee images with and without using side information. At $20\times$ acceleration, side information significantly aids reconstruction while its absence results in loss of fine anatomical details highlighted by the yellow arrows. Fig. 2 (right) shows axial FLAIR images of the brain. At $36\times$ total acceleration, side information markedly enhances the reconstruction fidelity, whereas its absence results in loss of essential features.

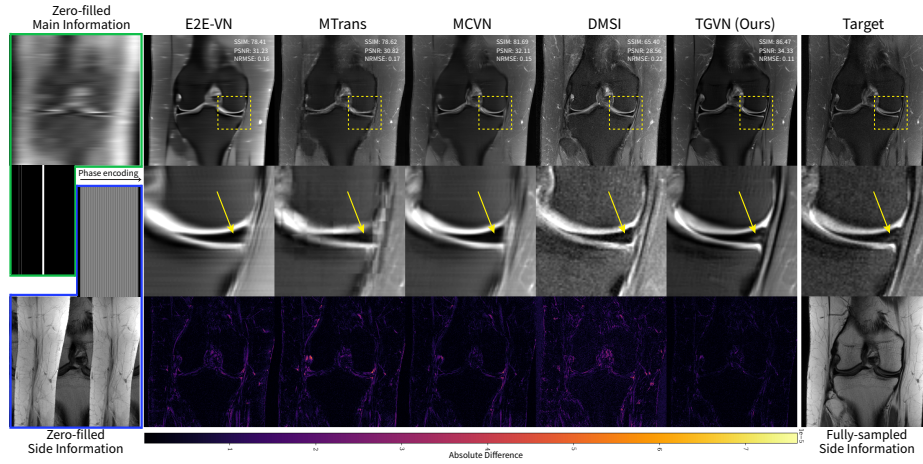


Fig. 3: **Knee image reconstructions from K1 demonstrating TGVN’s effective use of side information.** TGVN produces high-quality images at a challenging $20\times$ undersampling, outperforming multiple recently-proposed DL-based baselines. The *meniscus tear* (yellow arrow) is clearly visible only in the TGVN reconstruction, despite being obscured in the fully-sampled side information. **Top:** Full field-of-view images. **Middle:** Undersampling masks for main and side information (left), and zoomed-in regions (center, right). **Bottom:** Zero-filled reconstruction of undersampled side information used as input (left), absolute difference maps (center) between the target image and various reconstructions, using a consistent color map shown at bottom, and a fully sampled image of the side information (right, shown for illustration only. Models only had access to $2\times$ undersampled side information in K1). TGVN exhibits the smallest error and the best reconstruction metrics.

Fig. 3 compares TGVN reconstructions against several baselines using side information. MTrans and MCVN exhibit significant blurring of anatomical features, and DMSI suffers severely from noise amplification, which is seen clearly in the absolute difference images. The output of TGVN is significantly superior: both overall sharpness and assorted anatomical details are better preserved in the TGVN reconstructions. Furthermore, the meniscus tear highlighted by yellow arrows in the zoomed-in regions is distinctly more noticeable with TGVN. Notably, the meniscus tear is not well visualized in the image corresponding to the side information, which demonstrates that the TGVN is not relying excessively on features copied directly from the side information. Table 2 presents quantitative results, showing that TGVN achieves the best performance across all metrics with statistically significant improvements.

Overall, Fig. 2 and Table 2 address Q1 by demonstrating that relevant side information can significantly enhance reconstruction quality, while Fig. 3 and

Table 2 answer Q2 by showing that TGVN leverages side information more effectively and reliably than existing methods having comparable model capacity.

5 Conclusion

We introduced the **Trust-Guided Variational Network (TGVN)**, which demonstrates the power of leveraging side information in solving LIPs, with specific application to MR image reconstruction. Our key finding is that, when incorporated effectively, subject-specific side information can significantly improve reconstruction quality and preserve anatomical and pathological features while avoiding hallucinations, even at undersampling levels as high as $20\times$, and even with moderately-undersampled or low-quality side information. Note that TGVN is trained end-to-end, so side information cannot degrade the reconstruction quality. In the worst case (i.e., \mathbf{s} and \mathbf{x} are conditionally independent given \mathbf{k}), TGVN simply learns to ignore the side information. So far, we have used only complementary-contrast measurements from the same MRI exam as side information. However, a wide range of other side information can be encoded as needed then decoded into the image domain by the \mathcal{H} block. Essentially, \mathcal{H} maps the side information into the complex-valued, coil-combined MR image domain, making $\mathbf{x} - \mathcal{H}(\mathbf{s}; \gamma)$ well-defined in the trust-guidance term. In future work, we intend to incorporate prior images and text reports as side information.

Acknowledgments. This work was supported in part by the National Institute of Biomedical Imaging and Bioengineering (NIH P41 EB017183) and the National Science Foundation (NSF Award 1922658).

Disclosure of Interests. Not in connection with this work, Dr. Sodickson receives fees and holds stock options as a scientific advisor for Ezra, and receives royalties from a patent on deep learning-based image reconstruction licensed by Siemens Healthineers.

References

1. Aggarwal, H.K., Mani, M.P., Jacob, M.: MoDL: Model-Based Deep Learning Architecture for Inverse Problems. *IEEE Trans. Med. Imaging* **38**(2), 394–405 (2019)
2. Bilgic, B., Goyal, V.K., Adalsteinsson, E.: Multi-contrast reconstruction with Bayesian compressed sensing. *Magn. Reson. Med.* **66**(6), 1601–1615 (2011)
3. Bowsher, J.E., Johnson, V.E., Turkington, T.G., Jaszczak, R.J., Floyd, C., Coleman, R.E.: Bayesian reconstruction and use of anatomical a priori information for emission tomography. *IEEE Trans. Med. Imaging* **15**(5), 673–686 (1996)
4. Chen, J., Jiang, J., Wu, F., Zheng, J.: Null Space Matters: Range-Null Decomposition for Consistent Multi-Contrast MRI Reconstruction. In: *Proc. AAAI Conf. Artif. Intell.* pp. 1081–1090 (2024)
5. Dar, S.U.H., Yurt, M., Shahdloo, M., Ildiz, M.E., Tınaz, B., Çukur, T.: Prior-Guided Image Reconstruction for Accelerated Multi-Contrast MRI via Generative Adversarial Networks. *IEEE J. Sel. Topics Signal Process.* **14**(6), 1072–1087 (2020)

6. Ehrhardt, M.J., Betcke, M.M.: Multicontrast MRI Reconstruction with Structure-Guided Total Variation. *SIAM J. Imaging Sci.* **9**(3), 1084–1106 (2016)
7. Feng, C.M., Yan, Y., Chen, G., Xu, Y., Hu, Y., Shao, L., Fu, H.: Multimodal Transformer for Accelerated MR Imaging. *IEEE Trans. Med. Imaging* **42**(10), 2804–2816 (2022)
8. Fessler, J.A., Clinthorne, N.H., Rogers, W.L.: Regularized emission image reconstruction using imperfect side information. *IEEE Trans. Nucl. Sci.* **39**(5), 1464–1471 (1992)
9. Gindi, G., Lee, M., Rangarajan, A., Zubal, I.G.: Bayesian reconstruction of functional images using anatomical information as priors. *IEEE Trans. Med. Imaging* **12**(4), 670–680 (1993)
10. Hammernik, K., Klatzer, T., Kobler, E., Recht, M.P., Sodickson, D.K., Pock, T., Knoll, F.: Learning a variational network for reconstruction of accelerated MRI data. *Magn. Reson. Med.* **79**(6), 3055–3071 (2018)
11. Hestenes, M.R., Stiefel, E., et al.: Methods of Conjugate Gradients for Solving Linear Systems. *J. Res. Natl. Bur. Stand.* **49**(6), 409–436 (1952)
12. Huang, J., Chen, C., Axel, L.: Fast multi-contrast MRI reconstruction. *Magn. Reson. Imaging* **32**(10), 1344–1352 (2014)
13. Hyun, C.M., Kim, H.P., Lee, S.M., Lee, S., Seo, J.K.: Deep learning for undersampled MRI reconstruction. *Phys. Med. Biol.* **63**(13), 135007 (2018)
14. Kelkar, V.A., Anastasio, M.: Prior Image-Constrained Reconstruction using Style-Based Generative Models. In: *Proc. Int. Conf. Mach. Learn.* pp. 5367–5377 (2021)
15. Knoll, F., Zbontar, J., Sriram, A., Muckley, M.J., Bruno, M., Defazio, A., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., et al.: fastMRI: A Publicly Available Raw k-Space and DICOM Dataset of Knee Images for Accelerated MR Image Reconstruction Using Machine Learning. *Radiol. Artif. Intell.* **2**(1), e190007 (2020)
16. Landweber, L.: An Iteration Formula for Fredholm Integral Equations of the First Kind. *Amer. J. Math.* **73**(3), 615–624 (1951)
17. Lei, P., Fang, F., Zhang, G., Xu, M.: Deep Unfolding Convolutional Dictionary Model for Multi-Contrast MRI Super-resolution and Reconstruction. In: *Proc. Int. Joint Conf. Artif. Intell.* pp. 1008–1016 (2023)
18. Lei, P., Fang, F., Zhang, G., Zeng, T.: Decomposition-Based Variational Network for Multi-Contrast MRI Super-Resolution and Reconstruction. In: *Proc. IEEE Int. Conf. Comput. Vis.* pp. 21296–21306 (2023)
19. Levac, B., Jalal, A., Ramchandran, K., Tamir, J.I.: MRI Reconstruction with Side Information using Diffusion Models. In: *Proc. Asilomar Conf. Signals, Syst., Comput.* pp. 1436–1442 (2023)
20. Lustig, M., Donoho, D.L., Pauly, J.M.: Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.* **58**(6), 1182–1195 (2007)
21. Lustig, M., Donoho, D.L., Santos, J.M., Pauly, J.M.: Compressed Sensing MRI. *IEEE Signal Process. Mag.* **25**(2), 72–82 (2008)
22. Lyu, M., Mei, L., Huang, S., Liu, S., Li, Y., Yang, K., Liu, Y., Dong, Y., Dong, L., Wu, E.X.: M4Raw: A multi-contrast, multi-repetition, multi-channel MRI k-space dataset for low-field MRI research. *Sci. Data* **10**(1), 264 (2023)
23. Pruessmann, K.P., Weiger, M., Scheidegger, M.B., Boesiger, P.: SENSE: Sensitivity encoding for fast MRI. *Magn. Reson. Med.* **42**(5), 952–962 (1999)
24. Radmanesh, A., Muckley, M.J., Murrell, T., Lindsey, E., Sriram, A., Knoll, F., Sodickson, D.K., Lui, Y.W.: Exploring the Acceleration Limits of Deep Learning Variational Network-based Two-dimensional Brain MRI. *Radiol. Artif. Intell.* **4**(6), e210313 (2022)

25. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. pp. 234–241 (2015)
26. Schwab, J., Antholzer, S., Haltmeier, M.: Deep null space learning for inverse problems: convergence analysis and rates. *Inverse Problems* **35**(2), 025008 (2019)
27. Sodickson, D.K., Manning, W.J.: Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays. *Magn. Reson. Med.* **38**(4), 591–603 (1997)
28. Song, P., Weizman, L., Mota, J.F.C., Eldar, Y.C., Rodrigues, M.R.D.: Coupled Dictionary Learning for Multi-Contrast MRI Reconstruction. *IEEE Trans. Med. Imaging* **39**(3), 621–633 (2020)
29. Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C.L., Yakubova, N., Knoll, F., Johnson, P.: End-to-End Variational Networks for Accelerated MRI Reconstruction. In: Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. pp. 64–73 (2020)
30. Wang, Y., Hu, Y., Yu, J., Zhang, J.: GAN Prior based Null-Space Learning for Consistent Super-Resolution. In: Proc. AAAI Conf. Artif. Intell. pp. 2724–2732 (2023)
31. Wang, Y., Yu, J., Zhang, J.: Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model. In: Proc. Int. Conf. Learn. Represent. (2023)
32. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**(4), 600–612 (2004)
33. Xiang, J., Dong, Y., Yang, Y.: FISTA-Net: Learning a Fast Iterative Shrinkage Thresholding Network for Inverse Problems in Imaging. *IEEE Trans. Med. Imaging* **40**(5), 1329–1339 (2021)
34. Xiang, L., Chen, Y., Chang, W., Zhan, Y., Lin, W., Wang, Q., Shen, D.: Deep-Learning-Based Multi-Modal Fusion for Fast MR Reconstruction. *IEEE Trans. Biomed. Eng.* **66**(7), 2105–2114 (2019)
35. Zhao, R., Yaman, B., Zhang, Y., Stewart, R., Dixon, A., Knoll, F., Huang, Z., Lui, Y.W., Hansen, M.S., Lungren, M.P.: fastMRI+, Clinical pathology annotations for knee and brain fully sampled magnetic resonance imaging data. *Sci. Data* **9**(1), 152 (2022)
36. Zhou, B., Zhou, S.K.: DuDoRNet: Learning a Dual-Domain Recurrent Network for Fast MRI Reconstruction with Deep T1 Prior. In: Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. pp. 4273–4282 (2020)