

MRI IMAGE RECOVERY USING DAMPED DENOISING VECTOR AMP

Subrata Sarkar[†], Rizwan Ahmad^{*}, and Philip Schniter[†]

[†]Dept. ECE, The Ohio State Univ., Columbus, OH, 43210, {sarkar.51, schniter.1}@osu.edu

^{*}Dept. BME, The Ohio State Univ., Columbus, OH, 43210, rizwan.ahmad@osumc.edu

ABSTRACT

Motivated by image recovery in magnetic resonance imaging (MRI), we propose a new approach to solving linear inverse problems based on iteratively calling a deep neural-network, sometimes referred to as plug-and-play recovery. Our approach is based on the vector approximate message passing (VAMP) algorithm, which is known for mean-squared error (MSE)-optimal recovery under certain conditions. The forward operator in MRI, however, does not satisfy these conditions, and thus we design new damping and initialization schemes to help VAMP. The resulting DD-VAMP++ algorithm is shown to outperform existing algorithms in convergence speed and accuracy when recovering images from the fastMRI database for the practical case of Cartesian sampling.

1. INTRODUCTION

Magnetic resonance imaging (MRI) is a non-invasive diagnostic tool that provides excellent soft-tissue contrast without the use of ionizing radiation. The primary drawback of MRI is that the data acquisition process is inherently slow. Because the scan time scales with the number of measurements, there is a motivation to collect as few measurements as possible while still providing accurate image recovery. Although our work focuses on MRI, it has applicability to other inverse problems where the goal is to recover a richly structured signal from undersampled Fourier measurements.

In MRI, the Fourier-domain (or “k space”) measurements $\mathbf{y} \in \mathbb{C}^M$ can be modeled as

$$\mathbf{y} = \mathbf{A}\mathbf{x}_0 + \mathbf{w} \text{ with } \mathbf{A} = \mathbf{M}\mathbf{F}, \quad (1)$$

where $\mathbf{x}_0 \in \mathbb{C}^N$ is the rasterized N -pixel image that we aim to recover, $\mathbf{F} \in \mathbb{C}^{N \times N}$ is the unitary 2D discrete Fourier transform (DFT) matrix, $\mathbf{M} \in \mathbb{C}^{M \times N}$ is a sampling mask (containing rows of the identity matrix $\mathbf{I} \in \mathbb{R}^{N \times N}$), and $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}/\gamma_w)$ is circularly symmetric additive white Gaussian noise (AWGN) of precision γ_w (i.e., variance $1/\gamma_w$). Typical sampling masks are illustrated in Fig. 1. The ratio of pixels




Fig. 1: Typical sampling masks for acceleration $R = 4$

to measurements, $R \triangleq N/M$, is known as the “acceleration rate.” When $R > 1$, there are many $\mathbf{x}$ that lead to the same $\mathbf{y}$, and so it is essential to use prior knowledge of $\mathbf{x}$ in recovery.

Many methods have been proposed for MRI image recovery. As we describe in the next section, the classical approach is based on iterative optimization [1, 2]. Recently, it has been proposed to train deep neural networks (DNNs) to directly map compressed measurements $\mathbf{y}$ to image estimates $\hat{\mathbf{x}}$, e.g., [3, 4]. Although such DNNs are capable of excellent performance, training them requires huge fully-sampled k-space datasets (which may be difficult or impossible to obtain) and their performance can be degraded by changes in the sampling mask $\mathbf{M}$ between training and testing [5].

For these reasons, we focus on a hybrid approach where a DNN is iteratively called for *image denoising*. The DNN denoiser can be trained using relatively few images (via patches); no k-space data are needed. Furthermore, the training of the denoiser is divorced from the sampling mask. This hybrid approach is often referred to as “plug-and-play” recovery [5, 6]. Our approach is based on the vector AMP (VAMP) algorithm with two enhancements that lead to significant improvements in MRI image recovery: novel initialization and damping schemes. Experiments with fastMRI [4] knee images show advantages over existing plug-and-play approaches to MRI image recovery in speed and accuracy.

2. BACKGROUND

2.1. Optimization-based methods

The conventional approach to MRI image recovery [1, 2] is to pose and solve an optimization problem of the form

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \left\{ \frac{\gamma_w}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 + \phi(\mathbf{x}) \right\}, \quad (2)$$

This work was funded in part by the National Institutes of Health under grant R01HL135489 and by the National Science Foundation under grant CCF-1955587.

where the regularizer $\phi : \mathbb{C}^N \rightarrow \mathbb{R}_+$ aims to penalize $\mathbf{x}$ that are not valid images. It is common to choose $\phi(\mathbf{x}) = \lambda \|\Psi \mathbf{x}\|_1$ with a suitable (e.g., wavelet) transform Ψ and carefully chosen $\lambda > 0$. In this case, the regularization encourages sparsity in the wavelet coefficients $\Psi \mathbf{x}$.

Many algorithms have been designed to solve optimization problems of the form (2) with convex $\phi(\cdot)$ [2]. A famous one is the alternating directions method of multipliers (ADMM) [7], which for (2) iterates

$$\mathbf{x}^{t+1} = \arg \min_{\mathbf{x}} \left\{ \frac{\gamma_w}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 + \frac{\gamma}{2} \|\mathbf{x} - \mathbf{v}^t + \mathbf{u}^t\|^2 \right\} \quad (3a)$$

$$\mathbf{v}^{t+1} = \text{prox}_{\gamma^{-1}\phi}(\mathbf{x}^{t+1} + \mathbf{u}^t) \quad (3b)$$

$$\mathbf{u}^{t+1} = \mathbf{u}^t + (\mathbf{x}^{t+1} - \mathbf{v}^{t+1}), \quad (3c)$$

for $t = 0, 1, 2, \dots$, starting from $\mathbf{v}^0 = \mathbf{0} = \mathbf{u}^0$, where

$$\text{prox}_{\rho}(\mathbf{r}) \triangleq \arg \min_{\mathbf{x}} \left\{ \rho(\mathbf{x}) + \frac{1}{2} \|\mathbf{x} - \mathbf{r}\|^2 \right\}. \quad (4)$$

In (3), $\gamma > 0$ is a tunable stepsize that affects the speed of ADMM's convergence but not its fixed point.

2.2. Plug-and-play methods

It can be useful to interpret the optimization (2) as maximum a posteriori (MAP) estimation of $\mathbf{x}$. This can be seen using Bayes rule, $p(\mathbf{x}|\mathbf{y}) = p(\mathbf{y}|\mathbf{x})p(\mathbf{x})/p(\mathbf{y})$, as follows,

$$\hat{\mathbf{x}}_{\text{MAP}} \triangleq \arg \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}) \quad (5)$$

$$= \arg \min_{\mathbf{x}} \left\{ -\ln p(\mathbf{y}|\mathbf{x}) - \ln p(\mathbf{x}) \right\}, \quad (6)$$

and confirming that (6) matches (2) under the prior $p(\mathbf{x}) \propto e^{-\phi(\mathbf{x})}$ and likelihood $p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}; \mathbf{A}\mathbf{x}, \mathbf{I}/\gamma_w)$. Similarly, ADMM's prox step (3b) can be interpreted as MAP denoising, i.e., MAP estimation of $\mathbf{x} \sim p(\mathbf{x}) \propto e^{-\phi(\mathbf{x})}$ from the noise-corrupted measurement $\mathbf{r} = \mathbf{x} + \mathbf{w}$ with $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}/\gamma)$.

Leveraging this MAP denoising interpretation, Bouman et al. [6] proposed to replace $\text{prox}_{\gamma^{-1}\phi}$ in line (3b) with a call to a black-box image denoiser $\mathbf{f} : \mathbb{C}^N \rightarrow \mathbb{C}^N$ like BM3D [8] or DnCNN [9], and coined the approach “plug-and-play” (PnP) ADMM. In MRI, PnP methods tend to significantly outperform [5] the conventional optimization approach (2) to image recovery. Note, however, that when (3b) is replaced with a denoising step of the form $\mathbf{v}^{t+1} = \mathbf{f}(\mathbf{x}^{t+1} + \mathbf{u}^t)$, the stepsize γ *does* affect the fixed-point and thus must be tuned.

2.3. Approximate message passing

Approximate message passing (AMP) [10] is another iterative approach to computing the solution of (2) or (6), i.e., the MAP estimate of $\mathbf{x}$ from $\mathbf{y}$. But it can also [11] be used to approximate $\hat{\mathbf{x}}_{\text{MMSE}} \triangleq \mathbb{E}\{\mathbf{x}|\mathbf{y}\}$, the minimum mean-squared error (MMSE) estimate of $\mathbf{x}$ from $\mathbf{y}$. In general, AMP iterates

$$\mathbf{v}^{t+1} = \beta \cdot (\mathbf{y} - \mathbf{A}\mathbf{x}^t + \frac{1}{M} \mathbf{v}^t \text{tr}\{\nabla \mathbf{f}^t(\mathbf{x}^{t-1} + \mathbf{A}^H \mathbf{v}^t)\}) \quad (7a)$$

$$\tau^{t+1} = \frac{1}{M} \|\mathbf{v}^{t+1}\|^2 \quad (7b)$$

$$\mathbf{x}^{t+1} = \mathbf{f}^{t+1}(\mathbf{x}^t + \mathbf{A}^H \mathbf{v}^{t+1}) \quad (7c)$$

over $t = 0, 1, 2, \dots$, starting from $\mathbf{v}^0 = \mathbf{0} = \mathbf{x}^0$, where $\mathbf{f}^t(\cdot)$ is a Lipschitz denoising function, $\text{tr}\{\nabla \mathbf{f}^t(\mathbf{r})\}$ is the trace of the Jacobian of $\mathbf{f}^t$ at $\mathbf{r}$, and $\beta = N/\|\mathbf{A}\|_F^2$. When configured for MAP estimation, AMP uses the MAP denoiser $\mathbf{f}_{\text{MAP}}^t(\mathbf{r}) = \text{prox}_{\tau^t \phi}(\mathbf{r})$. When configured for MMSE estimation, AMP instead uses the MMSE denoiser $\mathbf{f}_{\text{MMSE}}^t(\mathbf{r}) = \mathbb{E}\{\mathbf{x}|\mathbf{r}\}$ for $\mathbf{r} = \mathbf{x} + \mathbf{w}$ with $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \tau^t \mathbf{I})$.

When the forward operator $\mathbf{A}$ is large and i.i.d. sub-Gaussian, the macroscopic behavior of AMP is rigorously characterized by a scalar state-evolution for any Lipschitz $\mathbf{f}^t$ [12, 13]. When $\mathbf{f}^t$ is the MMSE denoiser and the state-evolution has a unique fixed point, AMP provably converges to the MMSE-optimal estimate $\hat{\mathbf{x}}_{\text{MMSE}}$ [12, 13]. Although an exact MMSE denoiser for images is unknown, one could instead use a black-box approximation (like BM3D) within AMP, as proposed in [14] and called “denoising-AMP” (D-AMP). In that setting, the Jacobian term in (7a) is typically approximated using the approach from [15],

$$\text{tr}\{\nabla \mathbf{f}^t(\mathbf{r})\} \approx \epsilon^{-1} \mathbf{q}^H [\mathbf{f}^t(\mathbf{r} + \epsilon \mathbf{q}) - \mathbf{f}^t(\mathbf{r})], \quad (8)$$

for random $\mathbf{q} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and small positive ϵ .

For AMP to behave as expected, the forward operator $\mathbf{A}$ should resemble a typical realization of a large, i.i.d. sub-Gaussian matrix. To expand the applicability of AMP, a variation called “vector AMP (VAMP)” was proposed [16]. VAMP has essentially the same desirable properties of AMP (e.g., rigorous state evolution and provable MMSE denoising) but they hold for a much larger class of random matrices: right orthogonally invariant (ROI) ones [16, 17]. An ROI matrix has a singular value decomposition of the form $\mathbf{U}\mathbf{S}\mathbf{V}^H$, where $\mathbf{U}$ is unitary, $\mathbf{S}$ is diagonal, and $\mathbf{V}$ is Haar (i.e., uniformly distributed over the set of unitary matrices). A denoising VAMP (D-VAMP) was proposed in [18] that uses a black-box image denoiser $\mathbf{f}^t$ together with the Jacobian approximation (8).

2.4. AMP for MRI

While the PnP-ADMM algorithm requires manual tuning of both the ADMM parameter γ and the noise variance assumed by the denoiser $\mathbf{f}$, the AMP algorithm (7) is self-tuning; it estimates the denoiser-input variance τ^t at each iteration t . This motivates the application of AMP to MRI image recovery.

Unfortunately, neither AMP nor VAMP works as expected in MRI because $\mathbf{A}$ in (1) is neither i.i.d. nor ROI. In fact, MRI applications of AMP or VAMP often diverge. For this reason, several MRI-specific variations of AMP have been proposed. For example, [19] used D-AMP with $\beta = 1$ instead of the default value $\beta = N/\|\mathbf{A}\|_F^2 = N/M$, which acts to slow down the algorithm to help convergence but degrades the fixed points. A refinement based on a matrix-valued β was proposed in [20]. In [21], a hybrid between AMP and VAMP was proposed for the special case of point sampling (recall Fig. 1), called variable density (VD)-AMP. We numerically investigate these approaches in the sequel.

3. PROPOSED APPROACH

3.1. Damped Denoising VAMP

The proposed damped D-VAMP (DD-VAMP) algorithm is summarized in Alg. 1. Note that, in the absence of damping (i.e., $\theta = 1 = \zeta^t$), DD-VAMP reduces to D-VAMP from [18]. An explanation of each step is now provided.

Line 3 of Alg. 1 computes the MMSE/MAP estimate of $\mathbf{x}$ under the likelihood function $\mathcal{N}(\mathbf{y}; \mathbf{A}\mathbf{x}, \mathbf{I}/\gamma_w)$ and the pseudo-prior $\mathbf{x} \sim \mathcal{N}(\mathbf{x}; \mathbf{r}_2^t, \mathbf{I}/\gamma_2^t)$ via

$$\mathbf{g}(\mathbf{r}; \gamma) \triangleq \arg \min_{\mathbf{x}} \left\{ \frac{\gamma_w}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 + \frac{\gamma}{2} \|\mathbf{x} - \mathbf{r}\|^2 \right\} \quad (9)$$

$$= \mathbf{F}^H (\gamma_w \mathbf{M}^T \mathbf{M} + \gamma \mathbf{I})^{-1} (\gamma \mathbf{F} \mathbf{r} + \gamma_w \mathbf{M}^T \mathbf{y}), \quad (10)$$

where (9) holds for general $\mathbf{A}$ and (10) holds under (1). Because $\mathbf{M}$ is a masking matrix, the computational complexity of (10) is dominated by two FFTs. Line 4 computes the sensitivity of this linear estimator $\mathbf{g}$ as

$$\text{tr}\{\nabla \mathbf{g}(\mathbf{r}; \gamma)\}/N = \gamma \text{tr}\{(\gamma_w \mathbf{A}^H \mathbf{A} + \gamma \mathbf{I})^{-1}\}/N \quad (11)$$

$$= ((1 - \frac{\mathbf{M}}{N})\gamma_w + \gamma)/(\gamma_w + \gamma), \quad (12)$$

where (11) holds for general $\mathbf{A}$ and (12) holds under (1). Lines 5–6 perform “Onsager cancellation,” which produces the $(\mathbf{r}_1^t, \gamma_1^t)$ passed on to the denoising stage. If $\mathbf{f}^t$ is Lipschitz and $\mathbf{A}$ is large and ROI, then these latter terms obey the statistical model $\mathbf{r}_1^t = \mathbf{x}_0 + \mathcal{N}(\mathbf{0}, \mathbf{I}/\gamma_1^t)$ [17].

Lines 8–12 of Alg. 1 comprise the denoising stage, which calls the denoiser $\mathbf{f}^t$ and approximates its sensitivity $\text{tr}\{\nabla \mathbf{f}^t(\mathbf{r}_1^t)\}/N$ in a manner similar to D-AMP (see (8)). Note that the denoiser can leverage knowledge of the noise variance $\tau^t = 1/\gamma_1^t$. Lines 11–12 perform Onsager cancellation, producing $(\bar{\mathbf{r}}_2^{t+1}, \bar{\gamma}_2^{t+1})$ for use by the linear stage.

Before $(\bar{\mathbf{r}}_2^{t+1}, \bar{\gamma}_2^{t+1})$ are passed on, however, they are damped in lines 15–16 using the damping factor $\zeta^t \in (0, 1]$ chosen in line 14; with $\zeta^t = 1$ there would be no damping, and as ζ^t decreases there would be more damping. A similar form of damping is performed on α_1^t in line 10 using the factor θ . Intuitively, the goal of damping is to slow down VAMP without changing its fixed points. The $\mathbf{r}_2^t$ -damping in line 15 was shown to be sufficient to ensure the convergence of a double-loop version of VAMP in the strongly convex scenario [22]. But additional damping of γ_2^t and α_1^t is needed to stabilize D-VAMP for MRI image recovery.

Although γ_i -damping was suggested in [16, eq.(27)], the approach here is different in that damped quantities are first converted to amplitudes. That is, for a variance term like α_1 we damp the square-root, and for a precision term like γ_2 we damp the inverse square root. The α_1 -damping is new and motivated by the fact that the Jacobian approximation in line 9 can be very “noisy,” especially in the first iterations. Although it could be improved using more averaging, i.e.,

$$\text{tr}\{\nabla \mathbf{f}(\mathbf{r})\} \approx \frac{1}{K} \sum_{k=1}^K \epsilon^{-1} \mathbf{q}_k^H [\mathbf{f}(\mathbf{r} + \epsilon \mathbf{q}_k) - \mathbf{f}(\mathbf{r})], \quad (13)$$

Algorithm 1 Damped Denoising-VAMP (DD-VAMP)

Require: $\mathbf{r}_2^0 \in \mathbb{C}^N$, $\gamma_2^0 > 0$, $\theta \in (0, 1]$, $\mathbf{q} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

- 1: **for** $t = 0, \dots, T_{\max}$ **do**
 - 2: Linear Stage:
 - 3: $\mathbf{x}_2^t = \mathbf{g}(\mathbf{r}_2^t; \gamma_2^t)$
 - 4: $\alpha_2^t = \text{tr}\{\nabla \mathbf{g}(\mathbf{r}_2^t; \gamma_2^t)\}/N$
 - 5: $\mathbf{r}_1^t = (\mathbf{x}_2^t - \alpha_2^t \mathbf{r}_2^t)/(1 - \alpha_2^t)$
 - 6: $\gamma_1^t = \gamma_2^t(1 - \alpha_2^t)/\alpha_2^t$
 - 7: Denoising:
 - 8: $\mathbf{x}_1^t = \mathbf{f}^t(\mathbf{r}_1^t)$
 - 9: $\bar{\alpha}_1^t = \epsilon^{-1} \mathbf{q}^H [\mathbf{f}^t(\mathbf{r}_1^t + \epsilon \mathbf{q}) - \mathbf{f}^t(\mathbf{r}_1^t)]$
 - 10: $\alpha_1^t = \{\theta(\bar{\alpha}_1^t)^{\frac{1}{2}} + (1 - \theta)(\alpha_1^{t-1})^{\frac{1}{2}}\}^2$ (damping)
 - 11: $\bar{\mathbf{r}}_2^{t+1} = (\mathbf{x}_1^t - \alpha_1^t \mathbf{r}_1^t)/(1 - \alpha_1^t)$
 - 12: $\bar{\gamma}_2^{t+1} = \gamma_1^t(1 - \alpha_1^t)/\alpha_1^t$
 - 13: Damping:
 - 14: Choose damping factor $\zeta^t \in (0, 1]$
 - 15: $\mathbf{r}_2^{t+1} = \zeta^t \bar{\mathbf{r}}_2^{t+1} + (1 - \zeta^t) \mathbf{r}_2^t$
 - 16: $\gamma_2^{t+1} = \{\zeta^t (\bar{\gamma}_2^{t+1})^{-\frac{1}{2}} + (1 - \zeta^t) (\gamma_2^t)^{-\frac{1}{2}}\}^{-2}$
 - 17: **end for**
-

with i.i.d. $\{\mathbf{q}_k\}_{k=1}^K$ and large K , this would require $K + 1$ denoiser-calls per VAMP iteration and thus be too expensive.

There are several ways to choose the damping factor ζ^t in line 14. One is to use the rule

$$\zeta^t = \frac{2}{1 + \max\{\gamma_1^t/\bar{\gamma}_2^{t+1}, \bar{\gamma}_2^{t+1}/\gamma_1^t\}} = 2 \min\{\alpha_1^t, \alpha_2^t\}, \quad (14)$$

which is sufficient in the double-loop strongly convex scenario [22]. Another is to choose some fixed $\zeta^t = \zeta$ for all t .

3.2. ADMM-PR and DD-VAMP++

Empirically, we find that DD-VAMP converges to fixed points that are similar or better than those of PnP-ADMM. However, we find that it can converge somewhat slowly as a result of damping. To circumvent this issue, we propose a carefully initialized version called “DD-VAMP++.” The challenge with initializing DD-VAMP, i.e., choosing $(\mathbf{r}_2^0, \gamma_2^0)$, is that the two values must be consistent. That is, if $\mathbf{r}_2^0$ is very good (i.e., $\approx \mathbf{x}_0$) but γ_2^0 is not matched to $\mathbf{r}_2^0$, then DD-VAMP will initially move $\mathbf{r}_2^t$ away from $\mathbf{r}_2^0$ while adjusting γ_2^t , and finally return $\mathbf{r}_2^t$ to the neighborhood of $\mathbf{r}_2^0$.

The proposed initialization approach exploits the fact that VAMP is closely related to the Peaceman-Rachford (PR) variant [23] of ADMM, which is the iteration (3) with an additional $\mathbf{u}$ -update of the form (3c) between (3a) and (3b). In particular, when VAMP is run with fixed $\gamma_1 = \gamma_2 = \gamma$ for all iterations t , it reduces to ADMM-PR [22].

For DD-VAMP++, we first run PnP-ADMM-PR for T_{swi} iterations, starting from $\mathbf{u}^0 = \mathbf{0} = \mathbf{v}^0$ and a well-chosen value of γ , and then switch to DD-VAMP, simply by allowing the values of γ_1 and γ_2 to update according to Alg. 1.

DD-VAMP++'s two adjustable parameters, T_{swi} and γ , can be tuned using training data, as described in the next section.

4. NUMERICAL EXPERIMENTS

We now present numerical experiments that compare the proposed DD-VAMP and DD-VAMP++ algorithms to the existing PnP-ADMM [6], VD-AMP [21], and BM3D-AMP-MRI [19] algorithms for MRI image recovery. For completeness, we also test the PnP-ADMM-PR algorithm, which is run during the initialization phase of DD-VAMP++.

For the DD-VAMP, DD-VAMP++, PnP-ADMM, and PnP-ADMM-PR algorithms, we used the DNN-based denoiser DnCNN [9]. (The BM3D denoiser [8] yields similar image quality but runs significantly slower.) Our D-VAMP algorithms used the DnCNN implementation in the D-AMP toolbox <https://github.com/ricedsp/D-AMP-Toolbox>, which calls an instance of DnCNN trained with a noise variance close to the requested one, $\tau^t = 1/\gamma_1^t$. VD-AMP and BM3D-AMP-MRI were both implemented using the respective groups' code. To denoise, VD-AMP uses SURE-tuned wavelet thresholding while BM3D-AMP-MRI uses a modified BM3D. All experiments used mid-slice, non-fat-suppressed, 128×128 knee images from fastMRI [4]. We focused on the Cartesian mask (see Fig. 1) at acceleration $R = N/M = 4$. The variance of the AWGN w in (1) was adjusted to achieve an SNR, $\|Ax\|^2/\|w\|^2$, of 40 dB.

The PnP-ADMM, PnP-ADMM-PR, DD-VAMP, and DD-VAMP++ algorithms (all initialized at zero) have tuning parameters that must be optimized for best performance. For PnP-ADMM and PnP-ADMM-PR we tuned γ in (3), for DD-VAMP we tuned γ_2^0 , and for DD-VAMP++ we tuned γ and the switching time T_{swi} . In all cases, these parameters were tuned over a grid using the following procedure. For each hypothesized value, the algorithm was used to recover 30 training images, the recovery-NMSE was averaged over iterations $T_{\text{meas}} = 35$ to $T_{\text{max}} = 150$, and the median was taken over the training images (to reduce the effect of outlier images). Finally, the parameter value that minimized the median NMSE was selected.

Figure 2 shows $\text{NMSE} \triangleq \|\hat{x} - x_0\|^2/\|x_0\|^2$ versus iteration with a Cartesian mask, which is the mask most commonly used in clinical practice. The top subplot shows median NMSE over the 30 training images for the tuned version of each algorithm. There, DD-VAMP++ converged fastest and to the best NMSE. DD-VAMP achieved the 2nd-best NMSE at 150 iterations, but converged slowly. The two versions of PnP-ADMM showed similar behavior, although PnP-ADMM-PR was less smooth and its final NMSE was slightly worse. The bottom subplot of Fig. 2 shows median NMSE over the 19 remaining fastMRI mid-slice non-fat-suppressed images, which were used as a test set. Here too, we see that DD-VAMP++ converged fastest to the best NMSE, with PnP-ADMM and PnP-ADMM-PR close behind.




Fig. 2: Image recovery NMSE versus iteration.




Fig. 3: Image recovered after 150 iterations. Subplot titles identify algorithm, NMSE (dB), and SSIM.

After 150 iterations, DD-VAMP is still converging to what seems to be an excellent terminal NMSE. BM3D-AMP-MRI converged to a significantly worse NMSE; we conjecture that this was due to its heuristic β modification and the mismatch between Cartesian-masked A and typical i.i.d. sub-Gaussian matrices. Finally, VD-AMP (which was designed specifically for point masks) diverged; it is apparently not compatible with Cartesian masks. In experiments that we conducted with point masks (not shown here for reasons of space), VD-AMP converged very quickly to an NMSE > 10 dB worse than DD-VAMP++ and PnP-ADMM, but consistent with the results in [21]. Evidently, VD-AMP's wavelet-thresholding denoiser is no match for BM3D or DnCNN.

Figure 3 shows example image recoveries after 150 iterations, along with NMSE (dB) and SSIM values. DD-VAMP++ achieved the best recovery in NMSE and SSIM.

5. REFERENCES

- [1] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 72–82, Mar. 2008.
- [2] J. A. Fessler, "Optimization methods for magnetic resonance image reconstruction," *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 33–40, 2020.
- [3] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4509–4522, 2017.
- [4] J. Zbontar, F. Knoll, A. Sriram, M. J. Muckley, M. Bruno, A. Defazio, M. Parente *et al.*, "fastMRI: An open dataset and benchmarks for accelerated MRI," *arXiv:1811.08839*, 2018.
- [5] R. Ahmad, C. A. Bouman, G. T. Buzzard, S. Chan, S. Liu, E. T. Reehorst, and P. Schniter, "Plug and play methods for magnetic resonance imaging," *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 105–116, 2020.
- [6] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Global Conf. Signal Info. Process.*, 2013, pp. 945–948.
- [7] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
- [8] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, 2007.
- [9] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, 2017.
- [10] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing," *Proc. Nat. Acad. Sci.*, vol. 106, no. 45, pp. 18 914–18 919, Nov. 2009.
- [11] —, "Message passing algorithms for compressed sensing: I. Motivation and construction," in *Proc. Inform. Theory Workshop*, Cairo, Egypt, Jan. 2010, pp. 1–5.
- [12] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inform. Theory*, vol. 57, no. 2, pp. 764–785, Feb. 2011.
- [13] R. Berthier, A. Montanari, and P.-M. Nguyen, "State evolution for approximate message passing with non-separable functions," *Inform. Inference*, 2019.
- [14] C. A. Metzler, A. Maleki, and R. G. Baraniuk, "BM3D-AMP: A new image recovery algorithm based on BM3D denoising," in *Proc. IEEE Int. Conf. Image Process.*, 2015, pp. 3116–3120.
- [15] S. Ramani, T. Blu, and M. Unser, "Monte-Carlo SURE: A black-box optimization of regularization parameters for general denoising algorithms," *IEEE Trans. Image Process.*, vol. 17, no. 9, pp. 1540–1554, 2008.
- [16] S. Rangan, P. Schniter, and A. K. Fletcher, "Vector approximate message passing," *IEEE Trans. Inform. Theory*, pp. 6664–6684, 2019.
- [17] A. K. Fletcher, P. Pandit, S. Rangan, S. Sarkar, and P. Schniter, "Plug-in estimation in high-dimensional linear inverse problems: A rigorous analysis," in *Proc. Neural Inform. Process. Syst. Conf.*, 2018, pp. 7440–7449.
- [18] P. Schniter, S. Rangan, and A. K. Fletcher, "Denoising-based vector approximate message passing," in *Proc. Intl. Biomed. Astronom. Signal Process. (BASP) Frontiers Workshop*, 2017, p. 77.
- [19] E. M. Eksioğlu and A. K. Tanc, "Denoising AMP for MRI reconstruction: BM3D-AMP-MRI," *SIAM J. Imag. Sci.*, vol. 11, no. 3, pp. 2090–2109, 2018.
- [20] S. Sarkar, "Solving linear and bilinear inverse problems using approximate message passing methods," Ph.D. dissertation, Ohio State, 2020.
- [21] C. Millard, A. T. Hess, B. Mailhé, and J. Tanner, "Approximate message passing with a colored aliasing model for variable density Fourier sampled images," *arXiv:2003.02701*, 2020.
- [22] A. K. Fletcher, M. Sahraee-Ardakan, S. Rangan, and P. Schniter, "Expectation consistent approximate inference: Generalizations and convergence," *arXiv:1602.07795*, Feb. 2016.
- [23] D. Gabay, "Applications of the method of multipliers to variational inequalities," in *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems*, M. Fortin and R. Glowinski, Eds. Elsevier, 1983, pp. 299–331.