
Plug-and-Play Methods Provably Converge with Properly Trained Denoisers

Ernest K. Ryu¹ Jialin Liu¹ Sicheng Wang² Xiaohan Chen² Zhangyang Wang² Wotao Yin¹

Abstract

Plug-and-play (PnP) is a non-convex framework that integrates modern denoising priors, such as BM3D or deep learning-based denoisers, into ADMM or other proximal algorithms. An advantage of PnP is that one can use pre-trained denoisers when there is not sufficient data for end-to-end training. Although PnP has been recently studied extensively with great empirical success, theoretical analysis addressing even the most basic question of convergence has been insufficient. In this paper, we theoretically establish convergence of PnP-FBS and PnP-ADMM, without using diminishing stepsizes, under a certain Lipschitz condition on the denoisers. We then propose real spectral normalization, a technique for training deep learning-based denoisers to satisfy the proposed Lipschitz condition. Finally, we present experimental results validating the theory.

1. Introduction

Many modern image processing algorithms recover or denoise an image through the optimization problem

$$\underset{x \in \mathbb{R}^d}{\text{minimize}} \quad f(x) + \gamma g(x),$$

where the optimization variable $x \in \mathbb{R}^d$ represents the image, $f(x)$ measures data fidelity, $g(x)$ measures noisiness or complexity of the image, and $\gamma \geq 0$ is a parameter representing the relative importance between f and g . Total variation denoising, inpainting, and compressed sensing fall under this setup. *A priori* knowledge of the image, such as that the image should have small noise, is encoded in $g(x)$. So $g(x)$ is small if x has small noise or complexity. *A posteriori* knowledge of the image, such as noisy or partial

measurements of the image, is encoded in $f(x)$. So $f(x)$ is small if x agrees with the measurements.

First-order iterative methods are often used to solve such optimization problems, and ADMM is one such method:

$$\begin{aligned} x^{k+1} &= \underset{x \in \mathbb{R}^d}{\operatorname{argmin}} \left\{ \sigma^2 g(x) + (1/2) \|x - (y^k - u^k)\|^2 \right\} \\ y^{k+1} &= \underset{y \in \mathbb{R}^d}{\operatorname{argmin}} \left\{ \alpha f(y) + (1/2) \|y - (x^{k+1} + u^k)\|^2 \right\} \\ u^{k+1} &= u^k + x^{k+1} - y^{k+1} \end{aligned}$$

with $\sigma^2 = \alpha\gamma$. Given a function h on $\mathbb{R}^d$ and $\alpha > 0$, define the proximal operator of h as

$$\operatorname{Prox}_{\alpha h}(z) = \underset{x \in \mathbb{R}^d}{\operatorname{argmin}} \left\{ \alpha h(x) + (1/2) \|x - z\|^2 \right\},$$

which is well-defined if h is proper, closed, and convex. Now we can equivalently write ADMM as

$$\begin{aligned} x^{k+1} &= \operatorname{Prox}_{\sigma^2 g}(y^k - u^k) \\ y^{k+1} &= \operatorname{Prox}_{\alpha f}(x^{k+1} + u^k) \\ u^{k+1} &= u^k + x^{k+1} - y^{k+1}. \end{aligned}$$

We can interpret the subroutine $\operatorname{Prox}_{\sigma^2 g} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ as a denoiser, i.e.,

$$\operatorname{Prox}_{\sigma^2 g} : \text{noisy image} \mapsto \text{less noisy image}$$

(For example, if σ is the noise level and $g(x)$ is the total variation (TV) norm, then $\operatorname{Prox}_{\sigma^2 g}$ is the standard Rudin–Osher–Fatemi (ROF) model (Rudin et al., 1992).) We can think of $\operatorname{Prox}_{\alpha f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ as a mapping enforcing consistency with measured data, i.e.,

$$\operatorname{Prox}_{\alpha f} : \text{less consistent} \mapsto \text{more consistent with data}$$

More precisely speaking, for any $x \in \mathbb{R}^d$ we have

$$g(\operatorname{Prox}_{\sigma^2 g}(x)) \leq g(x), \quad f(\operatorname{Prox}_{\alpha f}(x)) \leq f(x).$$

However, some state-of-the-art image denoisers with great empirical performance do not originate from optimization problems. Such examples include non-local means (NLM) (Buades et al., 2005), Block-matching and 3D filtering (BM3D) (Dabov et al., 2007), and convolutional neural

¹Department of Mathematics, University of California, Los Angeles, USA ²Department of Computer Science and Engineering, Texas A&M University, USA. Correspondence to: Wotao Yin <wotaoyinmath.ucla.edu>.

networks (CNN) (Zhang et al., 2017a). Nevertheless, such a denoiser $H_\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^d$ still has the interpretation

$$H_\sigma : \text{noisy image} \mapsto \text{less noisy image}$$

where $\sigma \geq 0$ is a noise parameter. Larger values of σ correspond to more aggressive denoising.

Is it possible to use such denoisers for a broader range of imaging problems, even though we cannot directly set up an optimization problem? To address this question, (Venkatakrishnan et al., 2013) proposed Plug-and-Play ADMM (PnP-ADMM), which simply replaces the proximal operator $\text{Prox}_{\sigma^2 g}$ with the denoiser H_σ :

$$\begin{aligned} x^{k+1} &= H_\sigma(y^k - u^k) \\ y^{k+1} &= \text{Prox}_{\alpha f}(x^{k+1} + u^k) \\ u^{k+1} &= u^k + x^{k+1} - y^{k+1}. \end{aligned}$$

Surprisingly and remarkably, this ad-hoc method exhibited great empirical success, and spurred much follow-up work.

Contribution of this paper. The empirical success of Plug-and-Play (PnP) naturally leads us to ask theoretical questions: When does PnP converge and what denoisers can we use? Past theoretical analysis has been insufficient.

The main contribution of this work is the convergence analyses of PnP methods. We study two Plug-and-play methods, Plug-and-play forward-backward splitting (PNP-FBS) and PNP-ADMM. For the analysis, we assume the denoiser H_σ satisfies a certain Lipschitz condition, formally defined as Assumption (A). Roughly speaking, the condition corresponds to the denoiser H_σ being close to the identity map, which is reasonable when the denoising parameter σ is small. In particular, we do not assume that H_σ is nonexpansive or differentiable since most denoisers do not have such properties. Under the assumption, we show that the PnP methods are contractive iterations.

We then propose real spectral normalization (realSN), a technique based on (Miyato et al., 2018) for more accurately constraining deep learning-based denoisers in their training to satisfy the proposed Lipschitz condition. Finally, we present experimental results validating our theory. Code used for experiments is available at: https://github.com/uclaopt/Provable_Plug_and_Play/

1.1. Prior work

Plug-and-play: Practice. The first PnP method was the Plug-and-play ADMM proposed in (Venkatakrishnan et al., 2013). Since then, other schemes such as the primal-dual method (Heide et al., 2014; Meinhardt et al., 2017; Ono, 2017), ADMM with increasing penalty parameter (Brifman et al., 2016; Chan et al., 2017), generalized approximate message passing (Metzler et al., 2016), Newton it-

eration (Buzzard et al., 2018), Fast Iterative Shrinkage-Thresholding Algorithm (Kamilov et al., 2017; Sun et al., 2018b), (stochastic) forward-backward splitting (Sun et al., 2019; 2018a;b), and alternating minimization (Dong et al., 2018) have been combined with the PnP technique.

PnP method reported empirical success on a large variety of imaging applications: bright field electron tomography (Sreehari et al., 2016), camera image processing (Heide et al., 2014), compression-artifact reduction (Dar et al., 2016), compressive imaging (Teodoro et al., 2016), deblurring (Teodoro et al., 2016; Rond et al., 2016; Wang & Chan, 2017), electron microscopy (Sreehari et al., 2017), Gaussian denoising (Buzzard et al., 2018; Dong et al., 2018), nonlinear inverse scattering (Kamilov et al., 2017), Poisson denoising (Rond et al., 2016), single-photon imaging (Chan et al., 2017), super-resolution (Brifman et al., 2016; Sreehari et al., 2016; Chan et al., 2017), diffraction tomography (Sun et al., 2019), Fourier ptychographic microscopy (Sun et al., 2018b), low-dose CT imaging (Venkatakrishnan et al., 2013; He et al., 2018; Ye et al., 2018; Lyu et al., 2019), hyperspectral sharpening (Teodoro et al., 2017; 2019), inpainting (Chan, 2019; Tirer & Giryes, 2019), and superresolution (Dong et al., 2018).

A wide range of denoisers have been used for the PnP framework. BM3D has been used the most (Heide et al., 2014; Dar et al., 2016; Rond et al., 2016; Sreehari et al., 2016; Chan et al., 2017; Kamilov et al., 2017; Ono, 2017; Wang & Chan, 2017), but other denoisers such as sparse representation (Brifman et al., 2016), non-local means (Venkatakrishnan et al., 2013; Heide et al., 2014; Sreehari et al., 2016; 2017; Chan, 2019), Gaussian mixture model (Teodoro et al., 2016; 2017; Shi & Feng, 2018; Teodoro et al., 2019), Patch-based Wiener filtering (Venkatakrishnan et al., 2013), nuclear norm minimization (Kamilov et al., 2017), deep learning-based denoisers (Meinhardt et al., 2017; He et al., 2018; Ye et al., 2018; Tirer & Giryes, 2019) and deep projection model based on generative adversarial networks (Chang et al., 2017) have also been considered.

Plug-and-play: Theory. Compared to the empirical success, much less progress was made on the theoretical aspects of PnP optimization. (Chan et al., 2017) analyzed convergence with a bounded denoiser assumption, establishing convergence using an increasing penalty parameter. (Buzzard et al., 2018) provided an interpretation of fixed points via “consensus equilibrium”. (Sreehari et al., 2016; Sun et al., 2019; Teodoro et al., 2017; Chan, 2019; Teodoro et al., 2019) proved convergence of PNP-ADMM and PNP-FBS with the assumption that the denoiser is (averaged) nonexpansive by viewing the methods to be fixed-point iterations. The nonexpansiveness assumption is not met with most denoisers as is, but (Chan, 2019) proposed modifications to the non-local means and Gaussian mixture model denoisers, which make them into linear filters, to enforce

nonexpansiveness. (Dong et al., 2018) presented a proof that relies on the existence of a certain Lyapunov function that is monotonic under H_σ , which holds only for simple H_σ . (Tirer & Giryes, 2019) analyzed a variant of PnP, but did not establish local convergence since their key assumption is only expected to be satisfied “in early iterations”.

Other PnP-type methods. There are other lines of works that incorporate modern denoisers into model-based optimization methods. The plug-in idea with half quadratic splitting, as opposed to ADMM, was discussed (Zoran & Weiss, 2011) and this approach was carried out with deep learning-based denoisers in (Zhang et al., 2017b). (Danielyan et al., 2012; Egiazarian & Katkovnik, 2015) use the notion of Nash equilibrium to propose a scheme similar to PnP. (Danielyan et al., 2010) proposed an augmented Lagrangian method similar to PnP. (Romano et al., 2017; Reehorst & Schniter, 2019) presented Regularization by Denoising (RED), which uses the (nonconvex) regularizer $x^T(x - H_\sigma(x))$ given a denoiser H_σ , and use denoiser evaluations in its iterations. (Fletcher et al., 2018) applies the plug-in approach to vector approximate message passing. (Yang et al., 2016; Fan et al., 2019) replaced both the proximal operator enforcing data fidelity and the denoiser with two neural networks and performed end-to-end training. Broadly, there are more works that incorporate model-based optimization with deep learning (Chen et al., 2018; Liu et al., 2019).

Image denoising using deep learning. Deep learning-based denoising methods have become state-of-the-art. (Zhang et al., 2017a) proposed an effective denoising network called DnCNN, which adopted batch normalization (Ioffe & Szegedy, 2015) and ReLU (Krizhevsky et al., 2012) into the residual learning (He et al., 2016). Other representative deep denoising models include the deep convolutional encoder-decoder with symmetric skip connection (Mao et al., 2016), N^3 Net (Plötz & Roth, 2018), and MWCNN (Liu et al., 2018). The recent FFDNet (Zhang et al., 2018) handles spatially varying Gaussian noise.

Regularizing Lipschitz continuity. Lipschitz continuity and its variants have started to receive attention as a means for regularizing deep classifiers (Bartlett et al., 2017; Bansal et al., 2018; Oberman & Calder, 2018) and GANs (Miyato et al., 2018; Brock et al., 2019). Regularizing Lipschitz continuity stabilizes training, improves the final performance, and enhances robustness to adversarial attacks (Weng et al., 2018; Qian & Wegman, 2019). Specifically, (Miyato et al., 2018) proposed to normalize all weights to be of unit spectral norms to thereby constrain the Lipschitz constant of the overall network to be no more than one.

2. PNP-FBS/ADMM and their fixed points

We now present the PnP methods we investigate in this work. We quickly note that although PNP-FBS and PNP-ADMM

are distinct methods, they share the same fixed points by Remark 3.1 of (Meinhardt et al., 2017) and Proposition 3 of (Sun et al., 2019).

We call the method

$$x^{k+1} = H_\sigma(I - \alpha \nabla f)(x^k) \quad (\text{PNP-FBS})$$

for any $\alpha > 0$, plug-and-play forward-backward splitting (PNP-FBS) or plug-and-play proximal gradient method.

We interpret PNP-FBS as a fixed-point iteration, and we say x^* is a fixed point of PNP-FBS if

$$x^* = H_\sigma(I - \alpha \nabla f)(x^*).$$

Fixed points of PNP-FBS have a simple, albeit non-rigorous, interpretation. An image denoising algorithm must trade off the two goals of making the image agree with measurements and making the image less noisy. PNP-FBS applies $I - \alpha \nabla f$ and H_σ , each promoting such objectives, repeatedly in an alternating fashion. If PNP-FBS converges to a fixed point, we can expect the limit to represent a compromise.

We call the method

$$\begin{aligned} x^{k+1} &= H_\sigma(y^k - u^k) \\ y^{k+1} &= \text{Prox}_{\alpha f}(x^{k+1} + u^k) \\ u^{k+1} &= u^k + x^{k+1} - y^{k+1} \end{aligned} \quad (\text{PNP-ADMM})$$

for any $\alpha > 0$, plug-and-play alternating directions method of multipliers (PNP-ADMM). We interpret PNP-ADMM as a fixed-point iteration, and we say (x^*, u^*) is a fixed point of PNP-ADMM if

$$\begin{aligned} x^* &= H_\sigma(x^* - u^*) \\ x^* &= \text{Prox}_{\alpha f}(x^* + u^*). \end{aligned}$$

If we let $y^k = x^*$ and $u^k = u^*$ in (PNP-ADMM), then we get $x^{k+1} = y^{k+1} = x^*$ and $u^{k+1} = u^k = u^*$. We call the method

$$\begin{aligned} x^{k+1/2} &= \text{Prox}_{\alpha f}(z^k) \\ x^{k+1} &= H_\sigma(2x^{k+1/2} - z^k) \\ z^{k+1} &= z^k + x^{k+1} - x^{k+1/2} \end{aligned} \quad (\text{PNP-DRS})$$

plug-and-play Douglas–Rachford splitting (PNP-DRS). We interpret PNP-DRS as a fixed-point iteration, and we say z^* is a fixed point of PNP-DRS if

$$\begin{aligned} x^* &= \text{Prox}_{\alpha f}(z^*) \\ x^* &= H_\sigma(2x^* - z^*). \end{aligned}$$

PNP-ADMM and PNP-DRS are equivalent. Although this is not surprising as the equivalence between convex ADMM and DRS is well known, we show the steps establishing equivalence in the supplementary document.

We introduce PNP-DRS as an analytical tool for analyzing PNP-ADMM. It is straightforward to verify that PNP-DRS can be written as $z^{k+1} = T(z^k)$, where

$$T = \frac{1}{2}I + \frac{1}{2}(2H_\sigma - I)(2\text{Prox}_{\alpha f} - I).$$

We use this form to analyze the convergence of PNP-DRS and translate the result to PNP-ADMM.

3. Convergence via contraction

We now present conditions that ensure the PnP methods are contractive and thereby convergent.

If we assume $2H_\sigma - I$ is nonexpansive, standard tools of monotone operator theory tell us that PnP-ADMM converges. However, this assumption is too strong. Chan et al. presented a counter example demonstrating that $2H_\sigma - I$ is not nonexpansive for the NLM denoiser (Chan et al., 2017).

Rather, we assume $H_\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^d$ satisfies

$$\|(H_\sigma - I)(x) - (H_\sigma - I)(y)\|^2 \leq \varepsilon^2 \|x - y\|^2 \quad (\text{A})$$

for all $x, y \in \mathbb{R}^d$ for some $\varepsilon \geq 0$. Since σ controls the strength of the denoising, we can expect H_σ to be close to identity for small σ . If so, Assumption (A) is reasonable.

Under this assumption, we show that the PNP-FBS and PNP-DRS iterations are **contractive** in the sense that we can express the iterations as $x^{k+1} = T(x^k)$, where $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ satisfies

$$\|T(x) - T(y)\| \leq \delta \|x - y\|$$

for all $x, y \in \mathbb{R}^d$ for some $\delta < 1$. We call δ the contraction factor. If x^* satisfies $T(x^*) = x^*$, i.e., x^* is a fixed point, then $x^k \rightarrow x^*$ geometrically by the classical Banach contraction principle.

Theorem 1 (Convergence of PNP-FBS). *Assume H_σ satisfies assumption (A) for some $\varepsilon \geq 0$. Assume f is μ -strongly convex, f is differentiable, and ∇f is L -Lipschitz. Then*

$$T = H_\sigma(I - \alpha \nabla f)$$

satisfies

$$\|T(x) - T(y)\| \leq \max\{|1 - \alpha\mu|, |1 - \alpha L|\}(1 + \varepsilon)\|x - y\|$$

for all $x, y \in \mathbb{R}^d$. The coefficient is less than 1 if

$$\frac{1}{\mu(1 + 1/\varepsilon)} < \alpha < \frac{2}{L} - \frac{1}{L(1 + 1/\varepsilon)}.$$

Such an α exists if $\varepsilon < 2\mu/(L - \mu)$.

Theorem 2 (Convergence of PNP-DRS). *Assume H_σ satisfies assumption (A) for some $\varepsilon \geq 0$. Assume f is μ -strongly convex and differentiable. Then*

$$T = \frac{1}{2}I + \frac{1}{2}(2H_\sigma - I)(2\text{Prox}_{\alpha f} - I)$$

satisfies

$$\|T(x) - T(y)\| \leq \frac{1 + \varepsilon + \varepsilon\alpha\mu + 2\varepsilon^2\alpha\mu}{1 + \alpha\mu + 2\varepsilon\alpha\mu}\|x - y\|$$

for all $x, y \in \mathbb{R}^d$. The coefficient is less than 1 if

$$\frac{\varepsilon}{(1 + \varepsilon - 2\varepsilon^2)\mu} < \alpha, \quad \varepsilon < 1.$$

Corollary 3 (Convergence of PNP-ADMM). *Assume H_σ satisfies assumption (A) for some $\varepsilon \in [0, 1)$. Assume f is μ -strongly convex. Then PNP-ADMM converges for*

$$\frac{\varepsilon}{(1 + \varepsilon - 2\varepsilon^2)\mu} < \alpha.$$

Proof. This follows from Theorem 2 and the equivalence of PNP-DRS and PNP-ADMM. $\square$

For PNP-FBS, we assume f is μ -strongly convex and ∇f is L -Lipschitz. For PNP-DRS and PNP-ADMM, we assume f is μ -strongly convex. These are standard assumptions that are satisfied in application such as image denoising/deblurring and single photon imaging. Strong convexity, however, does exclude a few applications such as compressed sensing, sparse interpolation, and super-resolution.

PNP-FBS and PNP-ADMM are distinct methods for finding the same set of fixed points. Sometimes, PNP-FBS is easier to implement since it only requires the computation of ∇f rather than $\text{Prox}_{\alpha f}$. On the other hand, PNP-ADMM has better convergence properties as demonstrated theoretically by Theorems 1 and 2 and empirically by our experiments.

The proof of Theorem 2 relies on the notion of “negatively averaged” operators of (Giselsson, 2017). It is straightforward to modify Theorems 1 and 2 to establish local convergence when Assumption (A) holds locally. Theorem 2 can be generalized to the case when f is strongly convex but non-differentiable using the notion of subgradients.

Recently, (Fletcher et al., 2018) proved convergence of “plug-and-play” vector approximate message passing, a method similar to ADMM, assuming Lipschitz continuity of the denoiser. Although the method, the proof technique, and the notion of convergence are different from ours, the similarities are noteworthy.

4. Real spectral normalization: enforcing Assumption (A)

We now present real spectral normalization, a technique for training denoisers to satisfy Assumption (A) and connect the practical implementations to the theory of Section 3.

4.1. Deep learning denoisers: SimpleCNN and DnCNN

We use a deep denoising model called **DnCNN** (Zhang et al., 2017a), which learns the residual mapping with a 17-layer CNN and reports state-of-the-art results on natural image denoising. Given a noisy observation $y = x + e$, where x is the clean image and e is noise, the residual mapping R outputs the noise, i.e., $R(y) = e$ so that $y - R(y)$ is the clean recovery. Learning the residual mapping is a popular approach in deep learning-based image restoration.

We also construct a simple convolutional encoder-decoder model for denoising and call it **SimpleCNN**. SimpleCNN consists of 4 convolutional layers, with ReLU and mean-square-error (MSE) loss and does not utilize any pooling or (batch) normalization.

We remark that realSN and the theory of this work is applicable to other deep denoisers. We use SimpleCNN to show that realSN is applicable to any CNN denoiser.

4.2. Lipschitz constrained deep denoising

Denote the denoiser (SimpleCNN or DnCNN) as $H(y) = y - R(y)$, where y is the noisy input and R is the residual mapping, i.e., $R(y) = y - H(y) = (I - H)(y)$. Enforcing Assumption (A) is equivalent to constraining the Lipschitz constant of $R(y)$. We propose a variant of the spectral normalization (SN) (Miyato et al., 2018) for this.

Spectral normalization. (Miyato et al., 2018) proposed to normalize the spectral norm of each layer-wise weight (with ReLU non-linearity) to one. Provided that we use 1-Lipschitz nonlinearities (such as ReLU), the Lipschitz constant of a layer is upper-bounded by the spectral norm of its weight, and the Lipschitz constant of the full network is bounded by the product of spectral norms of all layers (Gouk et al., 2018). To avoid the prohibitive cost of singular value decomposition (SVD) every iteration, SN approximately computes the largest singular values of weights using a small number of power iterations.

Given the weight matrix $W_l \in \mathbb{R}^{m \times n}$ of the l -th layer, vectors $u_l \in \mathbb{R}^m, v_l \in \mathbb{R}^n$ are initialized randomly and maintained in the memory to estimate the leading first left and right singular vector of W_l respectively. During each forward pass of the network, SN is applied to all layers $1 \leq l \leq L$ following the two-step routine:

1. Apply one step of the power method to update u_l, v_l :

$$v_l \leftarrow W_l^T u_l / \|W_l^T u_l\|_2, \quad u_l \leftarrow W_l v_l / \|W_l v_l\|_2$$

2. Normalize W_l with the estimated spectral norm:

$$W_l \leftarrow W_l / \sigma(W_l), \text{ where } \sigma(W_l) = u_l^T W_l v_l$$

While the basic methodology of SN suits our goal, the SN in (Miyato et al., 2018) uses a convenient but inexact implementation for convolutional layers. A convolutional layer

is represented by a four-dimensional kernel K_l of shape $(C_{\text{out}}, C_{\text{in}}, h, w)$, where h, w are kernel's height and width. SN reshapes K_l into a two-dimensional matrix $\tilde{K}_l$ of shape $(C_{\text{out}}, C_{\text{in}} \times h \times w)$ and regards $\tilde{K}_l$ as the matrix W_l above. This relaxation **underestimates** the true spectral norm of the convolutional operator (Corollary 1 of (Tsuzuku et al., 2018)) given by

$$\sigma(K_l) = \max_{x \neq 0} \|K_l * x\|_2 / \|x\|_2,$$

where x is the input to the convolutional layer and $*$ is the convolutional operator. This issue is not hypothetical. When we trained SimpleCNN with SN, the spectral norms of the layers were 3.01, 2.96, 2.82, and 1.31, i.e., SN failed to control the Lipschitz constant below 1.

Real spectral normalization. We propose an improvement to SN for convolutional¹ layers, called the **real spectral normalization** (realSN), to more accurately constrain the network's Lipschitz constant and thereby enforce Assumption (A).

In realSN, we directly consider the convolutional linear operator $\mathcal{K}_l : \mathbb{R}^{C_{\text{in}} \times h \times w} \rightarrow \mathbb{R}^{C_{\text{out}} \times h \times w}$, where h, w are input's height and width, instead of reshaping the convolution kernel K_l into a matrix. The power iteration also requires the conjugate (transpose) operator $\mathcal{K}_l^*$. It can be shown that $\mathcal{K}_l^*$ is another convolutional operator with a kernel that is a rotated version of the forward convolutional kernel; the first two dimensions are permuted and the last two dimensions are rotated by 180 degrees (Liu et al., 2019). Instead of two vectors u_l, v_l as in SN, realSN maintains $U_l \in \mathbb{R}^{C_{\text{out}} \times h \times w}$ and $V_l \in \mathbb{R}^{C_{\text{in}} \times h \times w}$ to estimate the leading left and right singular vectors respectively. During each forward pass of the neural network, realSN conducts:

1. Apply one step of the power method with operator $\mathcal{K}_l$:

$$V_l \leftarrow \mathcal{K}_l^*(U_l) / \|\mathcal{K}_l^*(U_l)\|_2, \\ U_l \leftarrow \mathcal{K}_l(V_l) / \|\mathcal{K}_l(V_l)\|_2.$$

2. Normalize the convolutional kernel K_l with estimated spectral norm:

$$K_l \leftarrow K_l / \sigma(\mathcal{K}_l), \text{ where } \sigma(\mathcal{K}_l) = \langle U_l, \mathcal{K}_l(V_l) \rangle$$

By replacing $\sigma(\mathcal{K}_l)$ with $\sigma(K_l)/c_l$, realSN can constrain the Lipschitz constant to any upper bound $C = \prod_{l=1}^L c_l$. Using the highly efficient convolution computation in modern deep learning frameworks, realSN can be implemented simply and efficiently. Specifically, realSN introduces three additional one-sample convolution operations for each layer in each training step. When we used a batch size of 128, the extra computational cost of the additional operations is mild.

¹We use stride 1 and zero-pad with width 1 for convolutions.




Figure 1. Histograms for experimentally verifying Assumption (A). The x-axis represents values of $\|(I - H_\sigma)(x) - (I - H_\sigma)(y)\|/\|x - y\|$ and the y-axis represents the frequency. The vertical red bar corresponds to the maximum value.

4.3. Implementation details

We refer to SimpleCNN and DnCNN regularized by realSN as **RealSN-SimpleCNN** and **RealSN-DnCNN**, respectively. We train them in the setting of Gaussian denoising with known fixed noise levels $\sigma = 5, 15, 40$. We used $\sigma = 5, 15$ for CS-MRI and single photon imaging, and $\sigma = 40$ for Poisson denoising. The regularized denoisers are trained to have Lipschitz constant (no more than) 1. The training data consists of images from the BSD500 dataset, divided into 40×40 patches. The CNN weights were initialized in the same way as (Zhang et al., 2017a). We train all networks using the ADAM optimizer for 50 epochs, with a mini-batch size of 128. The learning rate was 10^{-3} in the first 25 epochs, then decreased to 10^{-4} . On an Nvidia GTX 1080 Ti, DnCNN took 4.08 hours and realSN-DnCNN took 5.17 hours to train, so the added cost of realSN is mild.

5. Poisson denoising: validating the theory

Consider the Poisson denoising problem, where given a true image $x_{\text{true}} \in \mathbb{R}^d$, we observe independent Poisson random variables $y_i \sim \text{Poisson}((x_{\text{true}})_i)$, so $y_i \in \mathbb{N}$, for $i = 1, \dots, d$. For details and motivation for this problem setup, see (Rond et al., 2016).

For the objective function $f(x)$, we use the negative log-likelihood given by $f(x) = \sum_{i=1}^d \ell(x_i; y_i)$, where

$$\ell(x; y) = \begin{cases} -y \log(x) + x & \text{for } y > 0, x > 0 \\ 0 & \text{for } y = 0, x \geq 0 \\ \infty & \text{otherwise.} \end{cases}$$

We can compute $\text{Prox}_{\alpha f}$ elementwise with

$$\text{Prox}_{\alpha f}(x) = (1/2) \left(x - \alpha + \sqrt{(x - \alpha)^2 + 4\alpha y} \right).$$

The gradient of f is given by $\partial f / \partial x_i = -y_i / x_i + 1$ for $x_i > 0$ for $i = 1, \dots, d$. We set $\partial f / \partial x_i = 0$ when $x_i = 0$, although, strictly speaking, $\partial f / \partial x_i$ is undefined when $y_i > 0$ and $x_i = 0$. This does not seem to cause any problems in the experiments. Since we force the denoisers to output nonnegative pixel values, PNP-FBS never needs to evaluate $\partial f / \partial x_i$ for negative x_i .

For H_σ , we choose BM3D, SimpleCNN with and without realSN, and DnCNN with and without realSN. Note that

these denoisers are designed or trained for the purpose of **Gaussian denoising**, and here we integrate them into the PnP frameworks for Poisson denoising. We scale the image so that the peak value of the image, the maximum mean of the Poisson random variables, is 1. The y -variable was initialized to the noisy image for PnP-FBS and PnP-ADMM, and the u -variable was initialized to 0 for PnP-ADMM. We use the test set of 13 images in (Chan et al., 2017).

Convergence. We first examine which denoisers satisfy Assumption (A) with small ε . In Figure 1, we run PnP iterations of Poisson denoising on a single image (flag of (Rond et al., 2016)) with different models, calculate $\|(I - H_\sigma)(x) - (I - H_\sigma)(y)\|/\|x - y\|$ between the iterates and the limit, and plot the histogram. The maximum value of a histogram, marked by a vertical red bar, lower-bounds the ε of Assumption (A). Remember that Corollary 3 requires $\varepsilon < 1$ to ensure convergence of PnP-ADMM. Figure 1(a) proves that BM3D violates this assumption. Figures 1(b) and 1(c) and Figures 1(d) and 1(e) respectively illustrate that RealSN indeed improves (reduces) ε for SimpleCNN and DnCNN.

Figure 2 experimentally validates Theorems 1 and 2, by examining the average (geometric mean) contraction factor (defined in Section 3) of PnP-FBS and ADMM² iterations over a range of step sizes. Figure 2 qualitatively shows that PnP-ADMM exhibits more stable convergence than PnP-FBS. Theorem 1 ensures PnP-FBS is a contraction when α is within an interval and Theorem 2 ensures PnP-ADMM is a contraction when α is large enough. We roughly observe this behavior for the denoisers trained with RealSN.

Table 1. Average PSNR performance (in dB) on Poisson denoising (peak = 1) on the testing set in (Chan et al., 2017)

	BM3D	RealSN-DnCNN	RealSN-SimpleCNN
PNP-ADMM	23.4617	23.5873	18.7890
PNP-FBS	18.5835	22.2154	22.7280

Empirical performance. Our theory only concerns convergence and says nothing about the recovery performance of the output the methods converge to. We empirically verify that the PnP methods with RealSN, for which we analyzed convergence, yield competitive denoising results.

²We compute the contraction factor of the equivalent PnP-DRS.




Figure 2. Average contraction factor of 500 iterations for the Poisson denoising experiment. The x-axis represents the value of α and y-axis represents the contraction factor. Although lower means faster convergence, a smoother curve means the method is easier to tune and has more stable convergence.

We fix $\alpha = 0.1$ for all denoisers in PnP-ADMM, and $\alpha = 0.0125$ in PnP-FBS. For deep learning-based denoisers, we choose $\sigma = 40/255$. For BM3D, we choose $\sigma = \sqrt{\gamma\alpha}$ as suggested in (Rond et al., 2016) and use $\gamma = 1$.

Table 1 compares the PnP methods with BM3D, RealSN-DnCNN, and RealSN-SimpleCNN plugged in. In both PnP methods, one of the two denoisers using RealSN, for which we have theory, outperforms BM3D. It is interesting to observe that the PnP performance does not necessarily hinge on the strength of the plugged in denoiser and that different PnP methods favor different denoisers. For example, RealSN-SimpleCNN surpasses the much more sophisticated RealSN-DnCNN under PnP-FBS. However, RealSN-DnCNN leads to better, and overall best, denoising performance when plugged into PnP-ADMM.

6. More applications

We now apply PnP on two imaging problems and show that RealSN improves the reconstruction of PnP.³

³Code for our experiments in Sections 5 and 6 is available at https://github.com/uclaopt/Provable_Plug_and_Play/

Single photon imaging. Consider single photon imaging with quanta image sensors (QIS) (Fossum, 2011; Chan & Lu, 2014; Elgandy & Chan, 2016) with the model

$$z = \mathbf{1}(y \geq 1), \quad y \sim \text{Poisson}(\alpha_{sg} G x_{\text{true}})$$

where $x_{\text{true}} \in \mathbb{R}^d$ is the underlying image, $G : \mathbb{R}^d \rightarrow \mathbb{R}^{dK}$ duplicates each pixel to K pixels, $\alpha_{sg} \in \mathbb{R}$ is sensor gain, K is the oversampling rate, $z \in \{0, 1\}^{dK}$ is the observed binary photons. We want to recover x_{true} from z . The likelihood function is

$$f(x) = \sum_{j=1}^n -K_j^0 \log(e^{-\alpha_{sg} x_j / K}) - K_j^1 \log(1 - e^{-\alpha_{sg} x_j / K}),$$

where $K_j^1 = \sum_{i=1}^K z_{(j-1)K+i}$ is the number of ones in the j -th unit pixel, $K_j^0 = \sum_{i=1}^K 1 - z_{(j-1)K+i}$ is the number of zeros in the j -th unit pixel. The gradient of $f(x)$ is given by $\partial f / \partial x_j = (\alpha_{sg} / K)(K_j^0 - K_j^1 / (e^{\alpha_{sg} x_j / K} - 1))$ and the proximal operator of f is given in (Chan & Lu, 2014).

We compare PnP-ADMM and PnP-FBS respectively with the denoisers BM3D, RealSN-DnCNN, and RealSN-SimpleCNN. We take $\alpha_{sg} = K = 8$. The y -variable was initialized to K^1 for PnP-FBS and PnP-ADMM, and the u -variable was initialized to 0 for PnP-ADMM. All deep denoisers used in this experiment were trained with fixed noise level $\sigma = 15$. We report the PSNRs achieved at the 50th iteration, the 100th iteration, and the best PSNR values achieved within the first 100 iterations.

Table 2 reports the average PSNR results on the 13 images used in (Chan et al., 2017). Experiments indicate that PnP-ADMM methods constantly yields higher PSNR than the PnP-FBS counterparts using the same denoiser. The best overall PSNR is achieved using PnP-ADMM with RealSN-DnCNN, which shows nearly 1dB improvement over the result obtained with BM3D. We also observe that deep denoisers with RealSN make PnP converges more stably.

Compressed sensing MRI. Magnetic resonance imaging (MRI) is a widely-used imaging technique with a slow data acquisition. Compressed sensing MRI (CS-MRI) accelerates MRI by acquiring less data through downsampling. PnP is useful in medical imaging as we do not have a large amount of data for end-to-end training: we train the denoiser H_σ on natural images, and then “plug” it into the PnP framework to be applied to medical images. CS-MRI is described mathematically as

$$y = \mathcal{F}_p x_{\text{true}} + \varepsilon_e,$$

where $x_{\text{true}} \in \mathbb{C}^d$ is the underlying image, $\mathcal{F}_p : \mathbb{C}^d \rightarrow \mathbb{C}^k$ is the linear measurement model, $y \in \mathbb{C}^k$ is the measured data, and $\varepsilon_e \sim N(0, \sigma_e I_k)$ is measurement noise. We want to recover x_{true} from y . The objective function is

$$f(x) = (1/2) \|y - \mathcal{F}_p x\|^2.$$

Table 2. Average PSNR (in dB) of single photon imaging task on the test set of (Chan et al., 2017)

PnP-ADMM, $\alpha = 0.01$				PnP-FBS, $\alpha = 0.005$			
Average PSNR	BM3D	RealSN-DnCNN	RealSN-SimpleCNN	Average PSNR	BM3D	RealSN-DnCNN	RealSN-SimpleCNN
Iteration 50	30.0034	31.0032	29.2154	Iteration 50	28.7933	27.9617	29.0062
Iteration 100	30.0014	31.0032	29.2151	Iteration 100	29.0510	27.9887	29.0517
Best Overall	30.0474	31.0431	29.2155	Best Overall	29.5327	28.4065	29.3563

 Table 3. CS-MRI results (30% sample with additive Gaussian noise $\sigma_e = 15$) in PSNR (dB).

Sampling approach		Random		Radial		Cartesian	
Image		Brain	Bust	Brain	Bust	Brain	Bust
Zero-filling		9.58	7.00	9.29	6.19	8.65	6.01
TV (Lustig et al., 2005)		16.92	15.31	15.61	14.22	12.77	11.72
RecRF (Yang et al., 2010)		16.98	15.37	16.04	14.65	12.78	11.75
BM3D-MRI (Eksioglu, 2016)		17.31	13.90	16.95	13.72	14.43	12.35
PnP-FBS	BM3D	19.09	16.36	18.10	15.67	14.37	12.99
	DnCNN	19.59	16.49	18.92	15.99	14.76	14.09
	RealSN-DnCNN	19.82	16.60	18.96	16.09	14.82	14.25
	SimpleCNN	15.58	12.19	15.06	12.02	12.78	10.80
	RealSN-SimpleCNN	17.65	14.98	16.52	14.26	13.02	11.49
PnP-ADMM	BM3D	19.61	17.23	18.94	16.70	14.91	13.98
	DnCNN	19.86	17.05	19.00	16.64	14.86	14.14
	RealSN-DnCNN	19.91	17.09	19.08	16.68	15.11	14.16
	SimpleCNN	16.68	12.56	16.83	13.47	13.03	11.17
	RealSN-SimpleCNN	17.77	14.89	17.00	14.47	12.73	11.88

The gradient of $f(x)$ is given in (Liu et al., 2016) and the proximal operator of $f(x)$ is given in (Eksioglu, 2016). We use BM3D, SimpleCNN and DnCNN, and their variants by RealSN for the PnP denoiser H_σ .

We take $\mathcal{F}_p$ as the Fourier k-domain subsampling (partial Fourier operator). We tested random, radial, and Cartesian sampling (Eksioglu, 2016) with a sampling rate of 30%. The noise level σ_e is taken as 15/255.

We compare PnP frameworks with zero-filling, total-variation (TV) (Lustig et al., 2005), RecRF (Yang et al., 2010), and BM3D-MRI (Eksioglu, 2016)⁴. The parameters are taken as follows. For TV, the regularization parameter λ is taken as the best one from $\{a \times 10^b, a \in \{1, 2, 5\}, b \in \{-5, -4, -3, -2, -1, 0, 1\}\}$. For RecRF, the two parameters λ, μ are both taken from the above sets and the best results are reported. For BM3D-MRI, we set the “final noise level (the noise level in the last iteration)” as $2\sigma_e$, which is suggested in their MATLAB library. For PnP methods with H_σ as BM3D, we set $\sigma = 2\sigma_e$, take $\alpha \in \{0.1, 0.2, 0.5, 1, 2, 5\}$ and report the best results.

⁴Some recent deep-learning based methods (Yang et al., 2016; Kulkarni et al., 2016; Metzler et al., 2017; Zhang & Ghanem, 2018) are not compared here because we assume we do not have enough medical images for training.

For PnP-ADMM with H_σ as deep denoisers, we take $\sigma = \sigma_e = 15/255$ and $\alpha = 2.0$ uniformly for all the cases. For PnP-FBS with H_σ as deep denoisers, we take $\sigma = \sigma_e/3 = 5/255$ and $\alpha = 0.4$ uniformly. All deep denoisers are trained on BSD500 (Martin et al., 2001), a natural image data set; no medical image is used in training. The y -variable was initialized to the zero-filled solution for PnP-FBS and PnP-ADMM, and the u -variable was initialized to 0 for PnP-ADMM. Table 3 reports our results on CS-MRI, from which we can confirm the effectiveness of PnP frameworks. Moreover, using RealSN-DnCNN seems to be the clear winner over all. We also observe that PnP-ADMM generally outperforms PnP-FBS when using the same denoiser, which supports Theorems 1 and 2.

7. Conclusion

In this work, we analyzed the convergence of PnP-FBS and PnP-ADMM under a Lipschitz assumption on the denoiser. We then presented real spectral normalization a technique to enforce the proposed Lipschitz condition in training deep learning-based denoisers. Finally, we validate the theory with experiments.

Acknowledgements

We thank Pontus Giselsson for the discussion on negatively averaged operators and Stanley Chan for the discussion on the difficulties in establishing convergence of PnP methods. This work was partially supported by National Key R&D Program of China 2017YFB02029, AFOSR MURI FA9550-18-1-0502, NSF DMS-1720237, ONR N0001417121, and NSF RI-1755701.

References

- Bailion, J. B., Bruck, R. E., and Reich, S. On the asymptotic behavior of nonexpansive mappings and semigroups in Banach spaces. *Houston Journal of Mathematics*, 4(1), 1978.
- Bansal, N., Chen, X., and Wang, Z. Can we gain more from orthogonality regularizations in training deep networks? In *Advances in Neural Information Processing Systems*, pp. 4266–4276, 2018.
- Bartlett, P. L., Foster, D. J., and Telgarsky, M. J. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pp. 6240–6249, 2017.
- Bauschke, H. H. and Combettes, P. L. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer New York, 2nd edition, 2017.
- Brifman, A., Romano, Y., and Elad, M. Turning a denoiser into a super-resolver using plug and play priors. *2016 IEEE International Conference on Image Processing*, pp. 1404–1408, 2016.
- Brock, A., Donahue, J., and Simonyan, K. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019.
- Buades, A., Coll, B., and Morel, J.-M. A non-local algorithm for image denoising. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2005.
- Buzzard, G. T., Chan, S. H., Sreehari, S., and Bouman, C. A. Plug-and-play unplugged: Optimization-free reconstruction using consensus equilibrium. *SIAM Journal on Imaging Sciences*, 11(3):2001–2020, 2018.
- Chan, S. H. Performance analysis of Plug-and-Play ADMM: A graph signal processing perspective. *IEEE Transactions on Computational Imaging*, 2019.
- Chan, S. H. and Lu, Y. M. Efficient image reconstruction for gigapixel quantum image sensors. In *Signal and Information Processing (GlobalSIP), 2014 IEEE Global Conference on*, pp. 312–316. IEEE, 2014.
- Chan, S. H., Wang, X., and Elgendy, O. A. Plug-and-play ADMM for image restoration: Fixed-point convergence and applications. *IEEE Transactions on Computational Imaging*, 3(1):84–98, 2017.
- Chang, J. R., Li, C.-L., Poczos, B., and Kumar, B. V. One network to solve them all—solving linear inverse problems using deep projection models. In *2017 IEEE International Conference on Computer Vision*, pp. 5889–5898. IEEE, 2017.
- Chen, X., Liu, J., Wang, Z., and Yin, W. Theoretical linear convergence of unfolded ista and its practical weights and thresholds. In *Advances in Neural Information Processing Systems*, pp. 9061–9071, 2018.
- Combettes, P. L. and Yamada, I. Compositions and convex combinations of averaged nonexpansive operators. *Journal of Mathematical Analysis and Applications*, 425(1): 55–70, 2015.
- Dabov, K., Foi, A., Katkovnik, V., and Egiazarian, K. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8): 2080–2095, 2007.
- Danielyan, A., Katkovnik, V., and Egiazarian, K. Image deblurring by augmented Lagrangian with BM3D frame prior. In *Workshop on Information Theoretic Methods in Science and Engineering*, pp. 16–18, 2010.
- Danielyan, A., Katkovnik, V., and Egiazarian, K. BM3D frames and variational image deblurring. *IEEE Transactions on Image Processing*, 21(4):1715–1728, 2012.
- Dar, Y., Bruckstein, A. M., Elad, M., and Giryes, R. Postprocessing of compressed images via sequential denoising. *IEEE Transactions on Image Processing*, 25(7):3044–3058, 2016.
- Dong, W., Wang, P., Yin, W., and Shi, G. Denoising prior driven deep neural network for image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- Egiazarian, K. and Katkovnik, V. Single image super-resolution via BM3D sparse coding. In *2015 23rd European Signal Processing Conference*, pp. 2849–2853, 2015.
- Eksioglu, E. M. Decoupled algorithm for mri reconstruction using nonlocal block matching model: BM3D-MRI. *Journal of Mathematical Imaging and Vision*, 56(3):430–440, 2016.
- Elgendy, O. A. and Chan, S. H. Image reconstruction and threshold design for quanta image sensors. In *2016 IEEE International Conference on Image Processing*, pp. 978–982. IEEE, 2016.

- Fan, K., Wei, Q., Wang, W., Chakraborty, A., and Heller, K. InverseNet: Solving inverse problems with splitting networks. *IEEE International Conference on Multimedia and Expo*, 2019.
- Fletcher, A. K., Pandit, P., Rangan, S., Sarkar, S., and Schniter, P. Plug-in estimation in high-dimensional linear inverse problems: A rigorous analysis. In *Advances in Neural Information Processing Systems 31*, pp. 7451–7460. 2018.
- Fossum, E. The quanta image sensor (QIS): concepts and challenges. In *Imaging Systems and Applications*. Optical Society of America, 2011.
- Giselsson, P. Tight global linear convergence rate bounds for Douglas–Rachford splitting. *Journal of Fixed Point Theory and Applications*, 19(4):2241–2270, 2017.
- Gouk, H., Frank, E., Pfahringer, B., and Cree, M. Regularisation of neural networks by enforcing Lipschitz continuity. *arXiv preprint arXiv:1804.04368*, 2018.
- He, J., Yang, Y., Wang, Y., Zeng, D., Bian, Z., Zhang, H., Sun, J., Xu, Z., and Ma, J. Optimizing a parameterized plug-and-play ADMM for iterative low-dose CT reconstruction. *IEEE Transactions on Medical Imaging*, pp. 1–13, 2018.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Heide, F., Steinberger, M., Tsai, Y.-T., Rouf, M., Pająk, D., Reddy, D., Gallo, O., Liu, J., Heidrich, W., Egiazarian, K., Kautz, J., and Pulli, K. FlexISP: A flexible camera image processing framework. *ACM Transactions on Graphics (Proceedings SIGGRAPH Asia 2014)*, 33(6), 2014.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pp. 448–456, 2015.
- Kamilov, U. S., Mansour, H., and Wohlberg, B. A plug-and-play priors approach for solving nonlinear imaging inverse problems. *IEEE Signal Processing Letters*, 24(12):1872–1876, 2017.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pp. 1097–1105, 2012.
- Kulkarni, K., Lohit, S., Turaga, P., Kerviche, R., and Ashok, A. Reconnet: Non-iterative reconstruction of images from compressively sensed measurements. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 449–458, 2016.
- Liu, J., Chen, X., Wang, Z., and Yin, W. ALISTA: Analytic Weights Are As Good As Learned Weights in LISTA. *ICLR*, 2019.
- Liu, P., Zhang, H., Zhang, K., Lin, L., and Zuo, W. Multi-level Wavelet-CNN for image restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 886–895, 2018.
- Liu, Y., Zhan, Z., Cai, J.-F., Guo, D., Chen, Z., and Qu, X. Projected iterative soft-thresholding algorithm for tight frames in compressed sensing magnetic resonance imaging. *IEEE transactions on medical imaging*, 35(9): 2130–2140, 2016.
- Lustig, M., Santos, J. M., Lee, J.-H., Donoho, D. L., and Pauly, J. M. Application of compressed sensing for rapid MR imaging. *SPARS,(Rennes, France)*, 2005.
- Lyu, Q., Ruan, D., Hoffman, J., Neph, R., McNitt-Gray, M., and Sheng, K. Iterative reconstruction for low dose ct using Plug-and-Play alternating direction method of multipliers (ADMM) framework. *Proceedings of SPIE*, 10949, 2019.
- Mao, X., Shen, C., and Yang, Y.-B. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In *Advances in Neural Information Processing Systems*, pp. 2802–2810, 2016.
- Martin, D., Fowlkes, C., Tal, D., and Malik, J. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings of the 8th International Conference on Computer Vision*, volume 2, pp. 416–423, July 2001.
- Meinhardt, T., Moeller, M., Hazirbas, C., and Cremers, D. Learning proximal operators: Using denoising networks for regularizing inverse imaging problems. In *2017 International Conference on Computer Vision*, pp. 1799–1808, 2017.
- Metzler, C., Mousavi, A., and Baraniuk, R. Learned d-amp: Principled neural network based compressive image recovery. In *Advances in Neural Information Processing Systems*, pp. 1772–1783, 2017.
- Metzler, C. A., Maleki, A., and Baraniuk, R. G. From denoising to compressed sensing. *IEEE Transactions on Information Theory*, 62(9):5117–5144, 2016.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.

- Moreau, J. J. Proximité et dualité dans un espace Hilbertien. *Bulletin de la Société Mathématique de France*, 93:273–299, 1965.
- Oberman, A. M. and Calder, J. Lipschitz regularized deep neural networks converge and generalize. *arXiv preprint arXiv:1808.09540*, 2018.
- Ogura, N. and Yamada, I. Non-strictly convex minimization over the fixed point set of an asymptotically shrinking nonexpansive mapping. *Numerical Functional Analysis and Optimization*, 23(1-2):113–137, 2002.
- Ono, S. Primal-dual plug-and-play image restoration. *IEEE Signal Processing Letters*, 24(8):1108–1112, 2017.
- Plötz, T. and Roth, S. Neural nearest neighbors networks. In *Advances in Neural Information Processing Systems*, pp. 1095–1106, 2018.
- Polyak, B. T. *Introduction to Optimization*. Optimization Software Inc., New York, 1987.
- Qian, H. and Wegman, M. N. L2-nonexpansive neural networks. In *International Conference on Learning Representations*, 2019.
- Reehorst, E. T. and Schniter, P. Regularization by denoising: Clarifications and new interpretations. *IEEE Transactions on Computational Imaging*, 5(1):52–67, 2019.
- Romano, Y., Elad, M., and Milanfar, P. The little engine that could: Regularization by denoising (RED). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844, 2017.
- Rond, A., Giryes, R., and Elad, M. Poisson inverse problems by the plug-and-play scheme. *Journal of Visual Communication and Image Representation*, 41:96–108, 2016.
- Rudin, L. I., Osher, S., and Fatemi, E. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992.
- Ryu, E. K. and Boyd, S. Primer on monotone operator methods. *Applied and Computational Mathematics*, 15: 3–43, 2016.
- Shi, M. and Feng, L. Plug-and-play prior based on Gaussian mixture model learning for image restoration in sensor network. *IEEE Access*, 6:78113–78122, 2018.
- Sreehari, S., Venkatakrishnan, S. V., Wohlberg, B., Buzard, G. T., Drummy, L. F., Simmons, J. P., and Bouman, C. A. Plug-and-play priors for bright field electron tomography and sparse interpolation. *IEEE Transactions on Computational Imaging*, 2(4):408–423, 2016.
- Sreehari, S., Venkatakrishnan, S. V., Bouman, K. L., Simmons, J. P., Drummy, L. F., and Bouman, C. A. Multi-resolution data fusion for super-resolution electron microscopy. *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017.
- Sun, Y., Wohlberg, B., and Kamilov, U. S. Plug-in stochastic gradient method. *arXiv preprint arXiv:1811.03659*, 2018a.
- Sun, Y., Xu, S., Li, Y., Tian, L., Wohlberg, B., and Kamilov, U. S. Regularized Fourier ptychography using an online plug-and-play algorithm. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2018b.
- Sun, Y., Wohlberg, B., and Kamilov, U. S. An online plug-and-play algorithm for regularized image reconstruction. *IEEE Transactions on Computational Imaging*, 2019.
- Taylor, A. B., Hendrickx, J. M., and Glineur, F. Exact worst-case convergence rates of the proximal gradient method for composite convex minimization. *Journal of Optimization Theory and Applications*, 2018.
- Teodoro, A. M., Bioucas-Dias, J. M., and Figueiredo, M. A. T. Image restoration and reconstruction using variable splitting and class-adapted image priors. *IEEE International Conference on Image Processing*, 2016.
- Teodoro, A. M., Bioucas-Dias, J. M., and Figueiredo, M. A. T. Scene-adapted plug-and-play algorithm with convergence guarantees. In *2017 IEEE 27th International Workshop on Machine Learning for Signal Processing*, pp. 1–6, 2017.
- Teodoro, A. M., Bioucas-Dias, J. M., and Figueiredo, M. A. T. A convergent image fusion algorithm using scene-adapted Gaussian-mixture-based denoising. *IEEE Transactions on Image Processing*, 28(1):451–463, 2019.
- Tirer, T. and Giryes, R. Image restoration by iterative denoising and backward projections. *IEEE Transactions on Image Processing*, 28(3):1220–1234, 2019.
- Tsuzuku, Y., Sato, I., and Sugiyama, M. Lipschitz-margin training: Scalable certification of perturbation invariance for deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 6541–6550, 2018.
- Venkatakrishnan, S. V., Bouman, C. A., and Wohlberg, B. Plug-and-play priors for model based reconstruction. *2013 IEEE Global Conference on Signal and Information Processing*, pp. 945–948, 2013.
- Wang, X. and Chan, S. H. Parameter-free Plug-and-Play ADMM for image restoration. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 1323–1327, 2017.

- Weng, T.-W., Zhang, H., Chen, P.-Y., Yi, J., Su, D., Gao, Y., Hsieh, C.-J., and Daniel, L. Evaluating the robustness of neural networks: An extreme value theory approach. In *International Conference on Learning Representations*, 2018.
- Yang, J., Zhang, Y., and Yin, W. A fast alternating direction method for TVL1-L2 signal reconstruction from partial Fourier data. *IEEE Journal of Selected Topics in Signal Processing*, 4(2):288–297, 2010.
- Yang, Y., Sun, J., Li, H., and Xu, Z. Deep ADMM-Net for compressive sensing MRI. In *Advances in Neural Information Processing Systems*, pp. 10–18, 2016.
- Ye, D. H., Srivastava, S., Thibault, J., Sauer, K., and Bouman, C. Deep residual learning for model-based iterative CT reconstruction using plug-and-play framework. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 6668–6672, 2018.
- Zhang, J. and Ghanem, B. ISTA-Net: Interpretable optimization-inspired deep network for image compressive sensing. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- Zhang, K., Zuo, W., Chen, Y., Meng, D., and Zhang, L. Beyond a Gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017a.
- Zhang, K., Zuo, W., Gu, S., and Zhang, L. Learning deep CNN denoiser prior for image restoration. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2808–2817, 2017b.
- Zhang, K., Zuo, W., and Zhang, L. FFDNet: Toward a fast and flexible solution for cnn based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018.
- Zoran, D. and Weiss, Y. From learning models of natural image patches to whole image restoration. In *2011 International Conference on Computer Vision*, pp. 479–486, 2011.