

Provable Convergence of Plug-and-Play Priors with MMSE denoisers

Xiaojian Xu, *Student Member, IEEE*, Yu Sun, *Student Member, IEEE*, Jiaming Liu, *Student Member, IEEE*, Brendt Wohlberg, *Senior Member, IEEE*, and Ulugbek S. Kamilov, *Member, IEEE*

Abstract—Plug-and-play priors (PnP) is a methodology for regularized image reconstruction that specifies the prior through an image denoiser. While PnP algorithms are well understood for denoisers performing *maximum a posteriori probability* (MAP) estimation, they have not been analyzed for the *minimum mean squared error* (MMSE) denoisers. This letter addresses this gap by establishing the first theoretical convergence result for the iterative shrinkage/thresholding algorithm (ISTA) variant of PnP for MMSE denoisers. We show that the iterates produced by PnP-ISTA with an MMSE denoiser converge to a stationary point of some global cost function. We validate our analysis on sparse signal recovery in compressive sensing by comparing two types of denoisers, namely the *exact* MMSE denoiser and the *approximate* MMSE denoiser obtained by training a deep neural net.

Index Terms—Plug-and-play priors, inverse problems, computational imaging, iterative shrinkage/thresholding algorithm (ISTA), MMSE estimation, deep learning.

I. INTRODUCTION

THE recovery of an unknown signal from its noisy measurements is fundamental in signal processing. It often arises in the context of linear inverse problems, where the goal is to recover $\mathbf{x} \in \mathbb{R}^n$ from its noisy measurements

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{e}. \quad (1)$$

Here, $\mathbf{H} \in \mathbb{R}^{m \times n}$ represents the response of the acquisition system and $\mathbf{e} \in \mathbb{R}^m$ models the measurements noise.

The solution of ill-posed inverse problems is commonly formulated as a *regularized inversion*, expressed as an optimization problem

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \quad \text{with} \quad f(\mathbf{x}) = g(\mathbf{x}) + h(\mathbf{x}), \quad (2)$$

where g is the data-fidelity term and h is the regularizer or prior. Proximal algorithms [1] are widely used for solving (2) when the regularizer is nonsmooth. For example, the *iterative*

shrinkage/thresholding algorithm (ISTA) [2]–[5] is a standard approach for solving (2) via the iterations

$$\mathbf{z}^t = \mathbf{x}^{t-1} - \gamma \nabla g(\mathbf{x}^{t-1}) \quad (3a)$$

$$\mathbf{x}^t = \text{prox}_{\gamma h}(\mathbf{z}^t), \quad (3b)$$

where $\gamma > 0$ is the step-size parameter. The second step of ISTA relies on the *proximal operator* defined as

$$\text{prox}_{\tau h}(\mathbf{z}) := \arg \min_{\mathbf{x} \in \mathbb{R}^n} \left\{ \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2 + \tau h(\mathbf{x}) \right\}, \quad (4)$$

where $\tau > 0$ controls the influence of h . The proximal operator can be interpreted as a *maximum a posteriori probability* (MAP) estimator for the AWGN denoising problem

$$\mathbf{z} = \mathbf{x} + \mathbf{n} \quad \text{where} \quad \mathbf{x} \sim p_{\mathbf{x}}, \quad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \tau \mathbf{I}), \quad (5)$$

by setting $h(\mathbf{x}) = -\log(p_{\mathbf{x}}(\mathbf{x}))$. This perspective inspired the development of the PnP methodology [6], [7], where the proximal operator is replaced by a more general denoiser $\mathbf{D}(\cdot)$, such as BM3D [8] or DnCNN [9]. For example, PnP-ISTA [10] can be summarized as

$$\mathbf{z}^t = \mathbf{x}^{t-1} - \gamma \nabla g(\mathbf{x}^{t-1}) \quad (6a)$$

$$\mathbf{x}^t = \mathbf{D}_{\sigma}(\mathbf{z}^t), \quad (6b)$$

where by analogy to τ in (4), we introduce the parameter $\sigma > 0$ for controlling the relative strength of the denoiser $\mathbf{D}_{\sigma}$.

PnP algorithms have been shown to achieve state-of-the-art performance in many imaging problems [11]–[20]. Recent work has also provided theoretical convergence guarantees for PnP algorithms under various assumptions on the data-fidelity term and the denoiser [7], [11], [15], [21]–[24]. However, PnP has *not* been investigated for denoisers performing *minimum mean squared error* (MMSE) estimation

$$\mathbf{D}_{\sigma}(\mathbf{z}) = \mathbb{E}[\mathbf{x}|\mathbf{z}] = \int_{\mathbb{R}^n} \mathbf{x} p_{\mathbf{x}|\mathbf{z}}(\mathbf{x}|\mathbf{z}) d\mathbf{x}. \quad (7)$$

MMSE denoisers are “optimal” with respect to widely used image-quality metrics such as signal-to-noise ratio (SNR). However, they are generally *not* nonexpansive [16] and their direct computation is often intractable in high-dimensions [25]. Insights into the performance of PnP for MMSE denoisers are valuable as many denoisers (pre-trained CNNs, NLM, BM3D) can be interpreted as *approximate* or *empirical* MMSE denoisers [26]. In this letter, we show that PnP-ISTA with an MMSE denoiser converges to a stationary point of a certain (possibly nonconvex) cost function. To the best of our knowledge, this explicit link between PnP-ISTA and MMSE estimation

Research presented in this article was supported by the NSF award CCF-1813910 and by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under project number 20200061DR. (Corresponding author: Ulugbek S. Kamilov.)

X. Xu and Y. Sun are with the Department of Computer Science & Engineering, Washington University in St. Louis, St. Louis, MO 63130.

J. Liu is with the Department of Electrical & Systems Engineering, Washington University in St. Louis, St. Louis, MO 63130.

B. Wohlberg is with Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545 USA.

U. S. Kamilov (email: kamilov@wustl.edu) is with the Department of Computer Science & Engineering and the Department of Electrical & Systems Engineering, Washington University in St. Louis, St. Louis, MO 63130.

is absent from the current literature on PnP methods. Our analysis builds on an elegant formulation by Gribonval [27] that establishes a direct link between MMSE estimation and regularized inversion. We validate our analysis on sparse signal recovery in compressive sensing by comparing PnP-ISTA with two types of denoisers—the *exact* MMSE denoiser and the *approximate* MMSE denoiser obtained by training DnCNN [9] to minimize the mean squared error (MSE). Our simulations show convergence of PnP-ISTA for both denoisers, highlight close agreement between their performance, and illustrate the limitation of using an AWGN denoiser as a prior within ISTA.

II. THEORETICAL ANALYSIS

Our analysis requires three assumptions that serve as sufficient conditions for establishing theoretical convergence.

Assumption 1. *The prior p_x is non-degenerate over $\mathbb{R}^n$.*

As a reminder, a probability distribution p_x is *degenerate* over $\mathbb{R}^n$, if it is supported on a space of lower dimensions than n . Assumption 1 is satisfied by a large number of probability distributions, including those promoting sparsity, such as the Bernoulli-Gaussian prior adopted in our numerical evaluation. Consider the image set of the MMSE denoiser $\mathcal{X} := \text{Im}(\text{D}_\sigma)$. Assumption 1 is required for establishing an explicit link between (7) and the following regularizer [27]

$$h_{\text{mmse}}(\mathbf{x}) := \begin{cases} -\frac{1}{2\gamma} \|\mathbf{x} - \text{D}_\sigma^{-1}(\mathbf{x})\|^2 + \frac{\sigma^2}{\gamma} h_\sigma(\text{D}_\sigma^{-1}(\mathbf{x})) & \text{for } \mathbf{x} \in \mathcal{X} \\ +\infty & \text{for } \mathbf{x} \notin \mathcal{X}, \end{cases} \quad (8)$$

where $\gamma > 0$ is the step-size, $\text{D}_\sigma^{-1} : \mathcal{X} \rightarrow \mathbb{R}^n$ is the inverse mapping, which is well defined and smooth over $\mathcal{X}$ (see Appendix A), and $h_\sigma(\cdot) := -\log(p_z(\cdot))$, where p_z is the probability distribution of the AWGN corrupted observation (5). As discussed in Appendix A, the function h_{mmse} is smooth for any $\mathbf{x} \in \mathcal{X}$, which is the consequence of the smoothness of both D_σ^{-1} and h_σ . Note that it is generally easier to evaluate the denoiser D_σ (either exactly or approximately) than explicitly differentiating the expression (8).

Assumption 2. *The function g is continuously differentiable and has a Lipschitz continuous gradient with constant $L > 0$.*

This is a standard assumption used extensively in the analysis of gradient-based algorithms (see [28], for example).

Assumption 3. *The function f has a finite infimum $f^* > -\infty$.*

This mild assumption ensures that the function f is bounded from below. We can now establish the following result.

Theorem 1. *Run PnP-ISTA with a denoiser (7) under Assumptions 1-3 using a fixed step-size $0 < \gamma \leq 1/L$. Then, the sequence $\{f(\mathbf{x}^t)\}_{t \geq 0}$ with h defined in (8) monotonically decreases and $\|\nabla f(\mathbf{x}^t)\| \rightarrow 0$ as $t \rightarrow \infty$.*

The proof is provided in Appendix B. Theorem 1 establishes convergence of PnP-ISTA with MMSE denoisers to a stationary point of the problem (2) where h is specified in (8). The proof relies on the *majorization-minimization (MM)* strategy widely used in the context of both convex and nonconvex




Fig. 1. Convergence of PnP-ISTA for exact and approximate MMSE denoisers. The latter corresponds to DnCNN trained to minimize MSE. Average normalized cost $f(\mathbf{x}^t)/f(\mathbf{x}^0)$ is plotted against the iteration number with the shaded areas representing the range of values attained over 100 experiments. Note the monotonic decrease of the cost function f as predicted by our analysis as well as the excellent agreement of both denoisers.




Fig. 2. Convergence of PnP-ISTA for exact and approximate MMSE denoisers. The latter corresponds to DnCNN trained to minimize MSE. Average SNR (dB) is plotted against the iteration number with the shaded areas representing the range of values attained over 100 experiments. The SNR behavior of LASSO, implemented using ISTA with the ℓ_1 -norm prior, is also provided for reference. We highlight the excellent agreement of both denoisers and their superior SNR performance compared to the ℓ_1 regularization.

optimization [29]–[33]. It is important to note that the theorem does *not* assume that g or h are convex, or that the denoiser is nonexpansive. The convexity of h_{mmse} is equivalent to the log-concavity of p_z [34], which does not hold for a wide variety of priors, such as mixtures of Gaussians [16]. On the other hand, when the function g is convex and p_z is log-concave, Theorem 1 implies that PnP-ISTA achieves the global minimum of f . The denoiser D_σ is a proximal operator of a proper, closed, and convex function h if and only if D_σ is monotone and nonexpansive [35]. Finally, note that h_{mmse} depends on both γ and σ , both of which influence the relative weighting between g and h . This is the consequence of h_{mmse} being specified by *reverse engineering* the MMSE denoiser D_σ , which leads to the explicit dependence of the regularizer on the problem parameters [27].

III. NUMERICAL EVALUATION

We illustrate PnP-ISTA with both *exact* and *approximate* MMSE denoisers on the problem of sparse signal recovery in compressive sensing [36], [37]. We emphasize that our aim here is *not* to argue that ISTA is a superior sparse recovery algorithm, or that the MMSE denoiser as a superior signal prior. Rather, we seek to gain new insights into the behavior of PnP-ISTA with MMSE priors in highly controlled setting.

As a model for the sparse vector $\mathbf{x} \in \mathbb{R}^n$ with $n = 4096$, we consider a widely used independent and identi-




Fig. 3. Illustration of the recovery performance of PnP-ISTA for exact and approximate MMSE denoisers. Average SNR (dB) is plotted against the measurement rate (m/n) with the shaded areas representing the range of values attained over 100 experiments. We also provide the performance of LASSO and GAMP, two widely used algorithms for sparse recovery in compressive sensing. The figure highlights the suboptimality of both variants of PnP-ISTA compared to GAMP, which stems from their assumption that errors in every ISTA iteration are AWGN. One can also observe the remarkable agreement between two variants of PnP-ISTA in all experiments.

cally distributed (i.i.d.) Bernoulli-Gaussian distribution. Each component of $\mathbf{x}$ is thus generated from the distribution $p_x(x) = \alpha\phi_{\sigma_x}(x) + (1 - \alpha)\delta(x)$, where δ is the Dirac delta function and ϕ_{σ_x} is the Gaussian probability density function with zero mean and $\sigma_x > 0$ standard deviation. The parameter $0 \leq \alpha \leq 1$ in p_x controls the sparsity of the signal, and we fix $\sigma_x^2 = 1/\alpha$. Since the distribution $p_z = (\phi_{\sigma} * p_x)$ is not log-concave, the Bernoulli-Gaussian prior leads to a nonconvex regularizer [27]. Additionally, it is known that the corresponding denoiser is generally expansive (see the discussion in Section III of [27]). The entries of $\mathbf{H} \in \mathbb{R}^{m \times n}$ are generated as i.i.d. Gaussian random variables $\mathcal{N}(0, 1/m)$. For each experiment, we additionally corrupt measurements with AWGN of variance σ_e^2 corresponding to an input SNR of 20 dB. Accordingly, the data fidelity term is set as least-squares $g(\mathbf{x}) = (1/2)\|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2$. All plots are obtained by averaging results over 100 random trials.

We consider two reference signal recovery algorithms extensively used in compressive sensing. The first is the standard *least absolute shrinkage and selection operator* (LASSO) [38], which computes (2) with an ℓ_1 -norm regularizer $h(\mathbf{x}) = \lambda\|\mathbf{x}\|_1$. The regularization parameter $\lambda > 0$ of LASSO is optimized for each experiment to maximize SNR. The second reference method is the MMSE variant of the *generalized approximate message passing* (GAMP) [39], which is known to be nearly optimal for sparse signal recovery in compressive sensing [40]. The parameters of GAMP are set to the actual statistical parameters ($\alpha, \sigma_x, \sigma_e$) of the problem. While the suboptimality of ISTA to GAMP for random measurement matrices is well known, our aim is to illustrate the relative performance of “optimal” ISTA with the MMSE denoiser D_{σ} .

Since $\mathbf{x}$ is a vector with i.i.d. elements, the exact MMSE denoiser D_{σ} can be evaluated as a sequence of scalar integrals. As an approximate MMSE denoiser, we use DnCNN with depth 4 [9]. To that end, we train 9 different networks for the removal of AWGN at noise levels in the range from 0.01 to 0.37. The training was conducted over 2000 random realizations of the signal $\mathbf{x} \sim p_x$ using the ℓ_2 -loss. For each experiment, we select the network achieving the highest SNR

value under the scaling technique from [41].

Theorem 1 establishes monotonic convergence of PnP-ISTA in terms of the cost function f . This is illustrated in Fig. 1 for the measurement rate $m/n = 0.8$. The average normalized cost $f(\mathbf{x}^t)/f(\mathbf{x}^0)$ is plotted against the iteration number for both exact and approximate MMSE denoisers. The shaded areas indicate the range of values taken over 100 random trials. Fig. 2 illustrates the convergence behaviour of PnP-ISTA in terms of SNR (dB) for identical experimental setting by additionally including the SNR performance of LASSO as a reference. First, note the monotonic convergence of $\{f(\mathbf{x}^t)\}_{t \geq 0}$ as predicted by our analysis. Second, note the excellent agreement between two variants of PnP-ISTA. This close agreement is encouraging as deep neural nets have been extensively used as practical strategies for regularizing large-scale imaging problems. For example, the convergence of PnP-ISTA for a modified DnCNN, trained using the real spectral normalization, has been shown in [23].

The underlying assumption in PnP-ISTA is that errors within every ISTA iteration can be modeled as AWGN, which is known to be false [42]. This makes both exact and approximate MMSE denoisers “suboptimal” when used within PnP-ISTA. Unlike ISTA, GAMP explicitly ensures AWGN errors in every iteration for *random measurement matrices*, making it a valid upper bound in our experimental setting. Fig. 3 illustrates the suboptimality of “optimal” ISTA for different measurement rates, highlighting the necessity of developing more accurate error models for PnP iterations [43].

IV. CONCLUSION

We provide several new insights into the widely used PnP methodology by considering “optimal” MMSE denoisers. First, we have analyzed the convergence of PnP-ISTA for MMSE denoisers. Our analysis reveals the convergence of the algorithm even when the data-fidelity term is *nonconvex* and denoiser is *expansive*. Second, our simulations on sparse signal recovery illustrate the potential of *approximate* MMSE denoisers—obtained by training deep neural nets—to match the performance of the *exact* MMSE denoiser. The latter is intractable for high-dimensional imaging problems, while the former has been extensively used in practice. Third, our simulations highlight the *suboptimality* of “optimal” ISTA with an MMSE denoiser, due to the assumption that error within ISTA iterations are Gaussian. We hypothesize that a similar phenomenon is present in the context of imaging inverse problems, which suggests the possibility of performance improvements by using more refined statistical models for characterizing errors within PnP algorithms [43].

APPENDIX

A. MMSE Denoising as Proximal Operator

The relationship between MMSE estimation and regularized inversion has been established by Gribonval [27], and has also been discussed in other contexts [25], [34]. Our contribution is to formally connect this relationship to PnP algorithms, leading to their new interpretation for MMSE denoisers.

It is well known that the estimator (7) can be compactly expressed using *Tweedie's formula* [44]

$$D_\sigma(z) = z - \sigma^2 \nabla h_\sigma(z) \quad \text{with} \quad h_\sigma(z) = -\log(p_z(z)), \quad (9)$$

which can be obtained by differentiating (7) using the expression for the probability distribution

$$p_z(z) = (p_x * \phi_\sigma)(z) = \int_{\mathbb{R}^n} \phi_\sigma(z - x) p_x(x) dx, \quad (10)$$

where $\phi_\sigma(x)$ is the Gaussian probability density function. Since ϕ_σ is infinitely differentiable, so are p_z and D_σ . By differentiating D_σ , one can show that the Jacobian of D_σ is positive definite (see Lemma 2 in [27])

$$JD_\sigma(z) = \mathbf{I} - \sigma^2 Hh_\sigma(z) \succ 0, \quad z \in \mathbb{R}^n, \quad (11)$$

where Hh_σ denotes the Hessian matrix of the function h_σ . Finally, Assumption 1 also implies that D_σ is a *one-to-one* mapping from $\mathbb{R}^n$ to $\mathcal{X} = \text{Im}(D_\sigma)$, which means that $D^{-1} : \mathcal{X} \rightarrow \mathbb{R}^n$ is well defined and also infinitely differentiable over $\mathcal{X}$ (see Lemma 1 in [27]). This directly implies that the regularizer h in (8) is also infinitely differentiable for any $x \in \mathcal{X}$.

We will now show that

$$D_\sigma(z) = \text{prox}_{\gamma h}(z) = \arg \min_{x \in \mathbb{R}^n} \left\{ \frac{1}{2} \|x - z\|^2 + \gamma h(x) \right\}$$

where h is the (possibly nonconvex) function defined in (8). We thus establish that $u^* = z$ is the global minimizer of

$$\varphi(u) := \frac{1}{2} \|D_\sigma(u) - z\|^2 + \gamma h(D_\sigma(u)), \quad u \in \mathbb{R}^n.$$

By using (8) and (9), we get

$$\begin{aligned} \varphi(u) &= \frac{1}{2} \|D_\sigma(u) - z\|^2 - \frac{1}{2} \|D_\sigma(u) - u\|^2 + \sigma^2 h_\sigma(u) \\ &= \frac{1}{2} \|D_\sigma(u) - z\|^2 - \frac{\sigma^4}{2} \|\nabla h_\sigma(u)\|^2 + \sigma^2 h_\sigma(u). \end{aligned}$$

The gradient of φ is then given by

$$\begin{aligned} \nabla \varphi(z) &= [JD_\sigma(u)] (D_\sigma(u) + \sigma^2 \nabla h_\sigma(u) - z) \\ &= [JD_\sigma(u)] (u - z), \end{aligned}$$

where we used (11) in the second line and (9) in the third line. Consider $q(\nu) = \varphi(z + \nu u)$ and its derivative

$$q'(\nu) = \nabla \varphi(z + \nu u)^\top u = \nu u^\top [JD_\sigma(z + \nu u)] u.$$

From the positive definiteness of the Jacobian (11), we have $q'(\nu) < 0$ and $q'(\nu) > 0$ for $\nu < 0$ and $\nu > 0$, respectively. This implies that $\nu = 0$ is the global minimizer of q . Since $u \in \mathbb{R}^n$ is an arbitrary vector, we have that φ has no stationary point beyond $u^* = z$ and that $\varphi(z) < \varphi(u)$ for any $u \neq z$.

B. Convergence Analysis

Prior work has analyzed the convergence of PnP algorithms for contractive, nonexpansive, or bounded denoisers [7], [11], [16], [22]–[24]. Our analysis extends the prior work on PnP by analyzing convergence for MMSE denoisers without any assumptions on convexity of g and h or on nonexpansiveness

of D_σ . We adopt the *majorization-minimization (MM) strategy* widely used in nonconvex optimization [29]–[33].

Consider the following approximation of f at $s \in \mathbb{R}^n$

$$\mu(x, s) = g(s) + \nabla g(s)^\top (x - s) + \frac{1}{2\gamma} \|x - s\|^2 + h(x)$$

Assumption 2 implies that for any $0 < \gamma \leq 1/L$, we have

$$\mu(x, s) \geq f(x) \quad \text{and} \quad \mu(s, s) = f(s), \quad x, s \in \mathbb{R}^n. \quad (12)$$

We express (6) in the MM format

$$x^t = \arg \min_{x \in \mathbb{R}^n} \mu(x, x^{t-1}) = D_\sigma(x^{t-1} - \gamma \nabla g(x^{t-1})), \quad (13)$$

where from Appendix A, we know that $D_\sigma = \text{prox}_{\gamma h}$. Therefore, from (12) and (13), we directly have that

$$f(x^t) \leq \mu(x^t, x^{t-1}) \leq \mu(x^{t-1}, x^{t-1}) = f(x^{t-1}).$$

From Assumption 3, we know that f is bounded from below; therefore, the *monotone convergence theorem* implies that the sequence $\{f(x^t)\}_{t \geq 0}$ converges.

Consider the residual function r between μ and f

$$\begin{aligned} r(x) &= \mu(x, s) - f(x) \\ &= g(s) + \nabla g(s)^\top (x - s) + \frac{1}{2\gamma} \|x - s\|^2 - g(x). \end{aligned}$$

The definition of r implies that $r(s) = 0$ and $\nabla r(s) = 0$. Additionally, we have for any $x, y \in \mathbb{R}^n$

$$\begin{aligned} \|\nabla r(x) - \nabla r(y)\| &\leq (1/\gamma) \|x - y\| + \|\nabla g(x) - \nabla g(y)\| \\ &\leq (1/\gamma + L) \|x - y\| \leq (2/\gamma) \|x - y\|, \end{aligned}$$

where we used $0 < \gamma \leq 1/L$. The last inequality implies that ∇r is Lipschitz continuous with constant $2/\gamma$.

Denote by f^* the infimum of f and by $r_t(x) = \mu(x, x^{t-1}) - f(x) \geq 0$, $x \in \mathbb{R}^n$, the residual at t . Then,

$$\begin{aligned} r_t(x^t) &= \mu(x^t, x^{t-1}) - f(x^t) \leq f(x^{t-1}) - f(x^t) \\ &\Rightarrow \sum_{t=1}^{\infty} r_t(x^t) \leq (f(x^0) - f^*), \end{aligned}$$

where we used the fact that $f^* \leq \lim_{t \rightarrow \infty} f(x^t)$. This implies that $r_t(x^t) \rightarrow 0$ as $t \rightarrow \infty$.

Since ∇r_t is $(2/\gamma)$ -Lipschitz continuous, we have that

$$\begin{aligned} u &:= x^t - \frac{\gamma}{2} \nabla r_t(x^t) \\ &\Rightarrow r_t(u) \leq r_t(x^t) - \frac{\gamma}{4} \|\nabla r_t(x^t)\|^2. \end{aligned}$$

Since $r_t(x) \geq 0$, for all $x \in \mathbb{R}^n$, we have

$$\|\nabla r_t(x^t)\|^2 \leq \frac{4}{\gamma} (r_t(x^t) - r_t(u)) \leq \frac{4}{\gamma} r_t(x^t) \rightarrow 0,$$

as $t \rightarrow \infty$.

Finally, consider the gradient of f at $x^t \in \mathcal{X} = \text{Im}(D_\sigma)$

$$\|\nabla f(x^t)\| = \|\nabla_x \mu(x^t, x^{t-1}) - \nabla r_t(x^t)\| = \|\nabla r_t(x^t)\| \rightarrow 0,$$

as $t \rightarrow \infty$, where we used the fact that x^t is the minimizer of $\mu(x, x^{t-1})$. This concludes the proof.

REFERENCES

- [1] N. Parikh and S. Boyd, "Proximal Algorithms," *Foundations and Trends in Optimization*, vol. 1, no. 3, pp. 123–231, 2014.
- [2] M. A. T. Figueiredo and R. D. Nowak, "An EM Algorithm for Wavelet-Based Image Restoration," *IEEE Trans. Image Process.*, vol. 12, no. 8, pp. 906–916, Aug. 2003.
- [3] J. Bect, L. Blanc-Feraud, G. Aubert, and A. Chambolle, "A ℓ_1 -unified variational framework for image restoration," in *Proc. ECCV*, Springer, Ed., vol. 3024, New York, 2004, pp. 1–13.
- [4] I. Daubechies, M. Defrise, and C. D. Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.*, vol. 57, no. 11, pp. 1413–1457, Nov. 2004.
- [5] A. Beck and M. Teboulle, "Fast Gradient-Based Algorithm for Constrained Total Variation Image Denoising and Deblurring Problems," *IEEE Trans. Image Process.*, vol. 18, no. 11, pp. 2419–2434, Nov. 2009.
- [6] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-Play Priors for Model Based Reconstruction," in *Proc. IEEE Global Conf. Signal Process. and Inf. Process. (GlobalSIP)*, Austin, TX, USA, Dec. 2013, pp. 945–948.
- [7] S. Sreehari, S. V. Venkatakrishnan, B. Wohlberg, G. T. Buzzard, L. F. Drummy, J. P. Simmons, and C. A. Bouman, "Plug-and-Play Priors for Bright Field Electron Tomography and Sparse Interpolation," *IEEE Trans. Comp. Imag.*, vol. 2, no. 4, pp. 408–423, Dec. 2016.
- [8] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Color Image Denoising via Sparse 3D Collaborative Filtering with Grouping Constraint in Luminance-Chrominance Space," in *Proc. IEEE Int. Conf. Image Proc. (ICIP 2017)*, San Antonio, TX, USA, 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, Jul. 2017.
- [10] U. S. Kamilov, H. Mansour, and B. Wohlberg, "A Plug-and-Play Priors Approach for Solving Nonlinear Inverse Problems," *IEEE Signal Process. Lett.*, vol. 24, no. 12, pp. 1872–1876, Dec. 2017.
- [11] S. H. Chan, X. Wang, and O. A. Elgendy, "Plug-and-Play ADMM for Image Restoration: Fixed-Point Convergence and Applications," *IEEE Trans. Comp. Imag.*, vol. 3, no. 1, pp. 84–98, Mar. 2017.
- [12] S. Ono, "Primal-Dual Plug-and-Play Image Restoration," *IEEE Signal Process. Lett.*, vol. 24, no. 8, pp. 1108–1112, 2017.
- [13] K. Zhang, W. Zuo, S. Gu, and L. Zhang, "Learning Deep CNN Denoiser Prior for Image Restoration," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [14] T. Meinhardt, M. Moeller, C. Hazirbas, and D. Cremers, "Learning Proximal Operators: Using Denoising Networks for Regularizing Inverse Imaging Problems," in *Proc. IEEE Int. Conf. Comp. Vis. (ICCV)*, Venice, Italy, Oct. 2017, pp. 1799–1808.
- [15] G. T. Buzzard, S. H. Chan, S. Sreehari, and C. A. Bouman, "Plug-and-Play Unplugged: Optimization free reconstruction using consensus equilibrium," *SIAM J. Imaging Sci.*, vol. 11, no. 3, pp. 2001–2020, 2018.
- [16] A. M. Teodoro, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "A Convergent Image Fusion Algorithm Using Scene-Adapted Gaussian-Mixture-Based Denoising," *IEEE Trans. Image Process.*, vol. 28, no. 1, pp. 451–463, Jan. 2019.
- [17] W. Dong, P. Wang, W. Yin, G. Shi, F. Wu, and X. Lu, "Denoising Prior Driven Deep Neural Network for Image Restoration," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 10, pp. 2305–2318, Oct. 2019, conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [18] S. H. Chan, "Performance Analysis of Plug-and-Play ADMM: A Graph Signal Processing Perspective," *IEEE Transactions on Computational Imaging*, vol. 5, no. 2, pp. 274–286, Jun. 2019, conference Name: IEEE Transactions on Computational Imaging.
- [19] Y. Sun, S. Xu, Y. Li, L. Tian, B. Wohlberg, and U. S. Kamilov, "Regularized Fourier Ptychography Using an Online Plug-and-play Algorithm," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Process.*, Brighton, UK, May 2019, pp. 7665–7669.
- [20] G. Song, Y. Sun, J. Liu, Z. Wang, and U. S. Kamilov, "A new recurrent plug-and-play prior based on the multiple self-similarity network," *IEEE Signal Process. Lett.*, vol. 27, pp. 451–455, 2020.
- [21] A. Teodoro, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "Scene-Adapted plug-and-play algorithm with convergence guarantees," in *Proc. IEEE Int. Workshop on Machine Learning for Signal Processing*, Tokyo, Japan, Sep. 2017.
- [22] Y. Sun, B. Wohlberg, and U. S. Kamilov, "An Online Plug-and-Play Algorithm for Regularized Image Reconstruction," *IEEE Trans. Comput. Imaging*, vol. 5, no. 3, pp. 395–408, Sep. 2019.
- [23] E. K. Ryu, J. Liu, S. Wnag, X. Chen, Z. Wang, and W. Yin, "Plug-and-Play Methods Provably Converge with Properly Trained Denoisers," in *Proc. 36th Int. Conf. Machine Learning (ICML)*, Long Beach, CA, USA, Jun. 2019.
- [24] R. G. Gavaskar and K. N. Chaudhury, "Plug-and-play ISTA converges with kernel denoisers," *IEEE Signal Process. Lett.*, vol. 27, pp. 610–614, 2020.
- [25] A. Kazerouni, U. S. Kamilov, E. Bostan, and M. Unser, "Bayesian Denoising: From MAP to MMSE Using Consistent Cycle Spinning," *IEEE Signal Process. Lett.*, vol. 20, no. 3, pp. 249–252, Mar. 2013.
- [26] A. Buades, B. Coll, and J. M. Morel, "Image Denoising Methods. A New Nonlocal Principle," *SIAM Rev.*, vol. 52, no. 1, pp. 113–147, 2010.
- [27] R. Gribonval, "Should Penalized Least Squares Regression be Interpreted as Maximum A Posteriori Estimation?" *IEEE Trans. Signal Process.*, vol. 59, no. 5, pp. 2405–2410, May 2011.
- [28] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer Academic Publishers, 2004.
- [29] A. P. Dempster, E. Hazan, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc. Ser. B*, vol. 39, pp. 1–38, 1977.
- [30] K. Lange, D. R. Hunter, and I. Yang, "Optimization transfer using surrogate objective functions," *J. Comput. Graph. Statist.*, vol. 9, pp. 1–20, 2000.
- [31] A. Beck and M. Teboulle, "Gradient-based algorithms with applications to signal recovery problems," in *Convex Optimization in Signal Processing and Communications*. Cambridge, 2009, ch. 2, pp. 42–88.
- [32] M. Razaviyayn, M. Hong, and Z.-Q. Luo, "A stochastic successive minimization method for nonsmooth nonconvex optimization with applications to transceiver design in wireless communication networks," *Math. Program.*, vol. 157, pp. 515–545, 2016.
- [33] J. Mairal, "Incremental Majorization-Minimization Optimization with Application to Large-Scale Machine Learning," *SIAM J. Optim.*, vol. 25, no. 2, pp. 829–855, Jan. 2015.
- [34] R. Gribonval and P. Machart, "Reconciling 'priors' & 'priors' without prejudice?" in *Proc. Advances in Neural Information Processing Systems 26*, Lake Tahoe, NV, USA, Dec. 2013, pp. 2193–2201.
- [35] P. L. Combettes and J.-C. Pesquet, "Proximal thresholding algorithm for minimization over orthonormal bases," *SIAM J. Optim.*, vol. 18, no. 4, pp. 1351–1376, 2007.
- [36] E. J. Candès, J. Romberg, and T. Tao, "Robust Uncertainty Principles: Exact Signal Reconstruction From Highly Incomplete Frequency Information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [37] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [38] R. Tibshirani, "Regression and Selection via the Lasso," *J. R. Stat. Soc. Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.
- [39] S. Rangan, "Generalized Approximate Message Passing for Estimation with Random Linear Mixing," in *Proc. IEEE Int. Symp. Information Theory*, St. Petersburg, Russia, Aug. 2011, pp. 2168–2172.
- [40] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová, "Probabilistic Reconstruction in Compressed Sensing: Algorithms, Phase Diagrams, and Threshold Achieving Matrices," *J. Stat. Mech.*, vol. 2012, p. P08009, Aug. 2012.
- [41] X. Xu, J. Liu, Y. Sun, B. Wohlberg, and U. S. Kamilov, "Boosting the Performance of Plug-and-Play Priors via Denoiser Scaling," 2020, arXiv:2002.11546.
- [42] 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, November 2009.
- [43] N. Eslahi and A. Foi, "Anisotropic spatiotemporal regularization in compressive video recovery by adaptively modeling the residual errors as correlated noise," in *IEEE Image, Video, and Multidimensional Signal Processing Workshop*, 2018.
- [44] H. Robbins, "An empirical Bayes approach to statistics," *Proc. Third Berkeley Symp. on Math. Statist. and Prob.*, Vol. 1 (Univ. of Calif. Press, 1956), pp. 157–163, 1956.