

---

# Stochastic Deep Restoration Priors for Imaging Inverse Problems

---

Yuyang Hu<sup>1</sup> Albert Peng<sup>1</sup> Weijie Gan<sup>1</sup> Peyman Milanfar<sup>2</sup> Mauricio Delbracio<sup>2</sup> Ulugbek S. Kamilov<sup>1</sup>

## Abstract

Deep neural networks trained as image denoisers are widely used as priors for solving imaging inverse problems. We introduce *Stochastic deep Restoration Priors (ShaRP)*, a novel framework that stochastically leverages an ensemble of deep restoration models beyond denoisers to regularize inverse problems. By using generalized restoration models trained on a broad range of degradations beyond simple Gaussian noise, ShaRP effectively addresses structured artifacts and enables self-supervised training without fully sampled data. We prove that ShaRP minimizes an objective function involving a regularizer derived from the score functions of minimum mean square error (MMSE) restoration operators. We also provide theoretical guarantees for learning restoration operators from incomplete measurements. ShaRP achieves state-of-the-art performance on tasks such as magnetic resonance imaging reconstruction and single-image super-resolution, surpassing both denoiser- and diffusion-model-based methods without requiring retraining.

## 1. Introduction

Many problems in computational imaging, biomedical imaging, and computer vision can be viewed as *inverse problems*, where the goal is to recover an unknown image from its noisy and incomplete measurements. Inverse problems are typically ill-posed, thus requiring additional prior information for accurate image reconstruction. While many approaches have been proposed for implementing image priors, the current research focuses on methods based on deep learning (DL) (McCann et al., 2017; Ongie et al., 2020; Kamilov et al., 2023; Wen et al., 2023).

Deep neural networks trained as image denoisers are widely used for specifying image priors for solving *general* inverse

problems (Romano et al., 2017; Kadkhodaie & Simoncelli, 2021; Zhang et al., 2022). The combination of pre-trained Gaussian denoisers with measurement models has been shown to be effective in many inverse problems, including image super-resolution, deblurring, and medical imaging (see the recent reviews (Ahmad et al., 2020; Kamilov et al., 2023; Milanfar & Delbracio, 2024; Daras et al., 2024)). This success has led to active research on novel methods based on denoiser priors, their theoretical analyses, statistical interpretations, as well as connections to related approaches such as score matching and diffusion models (Venkatakrisnan et al., 2013; Chan et al., 2017; Romano et al., 2017; Buzzard et al., 2018; Reehorst & Schniter, 2019; Sun et al., 2019; Sun et al., 2019; Ryu et al., 2019; Cohen et al., 2021; Hurault et al., 2022b; Laumont et al., 2022; Gan et al., 2023a; Renaud et al., 2024; Xiao et al., 2024; Nathan et al., 2024).

Although priors based on Gaussian denoising models have been extensively studied, there is little research on leveraging priors from pre-trained restoration models that extend beyond Gaussian denoisers. In this paper, we present evidence that priors derived from deep models pre-trained as general restoration operators can surpass those trained exclusively for Gaussian denoising. We introduce a novel framework called *Stochastic deep Restoration Priors (ShaRP)*, which provides a principled approach to integrate an ensemble of general restoration models as priors to regularize inverse problems. By using more versatile restoration models, ShaRP improves upon traditional methods using Gaussian denoiser priors in two key ways: (a) ShaRP improved performance by using restoration models that are better suited for mitigating non-Gaussian structured artifacts arising during inference. (b) The restoration models in ShaRP can sometimes be directly trained in a self-supervised manner without fully-sampled measurement data.

We present novel theoretical and numerical results highlighting the potential of using an ensemble of restoration models as image priors. Our theoretical result introduces a novel notion of regularization for inverse problems corresponding to the average of likelihoods associated with the degraded observations of an image. The proposed regularizer has an intuitive interpretation as promoting solutions whose *multiple* degraded observations resemble realistic degraded images. We show that ShaRP seeks to minimize an objective function containing this regularizer. Our second theoretic

---

<sup>1</sup>WashU <sup>2</sup>Google. Correspondence to: Ulugbek S. Kamilov <kamilov@wustl.edu>.

cal result analyzes the possibility of learning the minimum mean squared error (MMSE) restoration operators directly from noisy and undersampled measurements. We numerically show the practical relevance of ShaRP by applying it to MRI reconstruction with varying undersampling patterns and rates, using a fixed-rate pre-trained MRI reconstruction network as a prior. We also show that ShaRP can use a pre-trained image deblurring model to perform single image super-resolution (SISR). Our numerical experiments show that ShaRP adapts the pre-trained restoration model as a prior, outperforming existing methods based on image denoisers and diffusion models, and achieving state-of-the-art results. Our experiments additionally highlight the benefit of using restoration models as priors by considering a setting where only undersampled and noisy MRI data is available for pre-training the prior. In such cases, self-supervised training of a restoration model is feasible, whereas training a Gaussian denoiser requires fully sampled data.

## 2. Background

**Inverse Problems.** Many computational imaging tasks can be formulated as inverse problems, where the goal is to reconstruct an unknown image  $\mathbf{x} \in \mathbb{R}^n$  from its corrupted measurement

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}, \quad (1)$$

where  $\mathbf{A} \in \mathbb{R}^{m \times n}$  is a measurement operator and  $\mathbf{e} \in \mathbb{R}^m$  is the noise. A common approach to addressing inverse problems is to formulate them as an optimization problem

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

where  $g$  is the data-fidelity term that quantifies the fit to the measurement  $\mathbf{y}$  and  $h$  is a regularizer that incorporates prior information on  $\mathbf{x}$ . For instance, typical functions used in imaging inverse problems are the least-squares term  $g(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2$  and the total variation (TV) regularizer  $h(\mathbf{x}) = \tau \|\mathbf{D}\mathbf{x}\|_1$ , where  $\mathbf{D}$  is the image gradient and  $\tau > 0$  is a regularization parameter.

**Deep Learning.** DL has emerged as a powerful tool for addressing inverse problems (McCann et al., 2017; Ongie et al., 2020; Wen et al., 2023). Instead of explicitly defining a regularizer, DL methods use deep neural networks (DNNs) to map the measurements to the desired images (Wang et al., 2016; Jin et al., 2017; Kang et al., 2017; Chen et al., 2017; Delbracio et al., 2021; Delbracio & Milanfar, 2023). Model-based DL (MBDL) is a widely-used sub-family of DL algorithms that integrate physical measurement models with priors specified using CNNs (see reviews by (Ongie et al., 2020; Monga et al., 2021)). The literature of MBDL is vast, but some well-known examples include plug-and-play priors (PnP), regularization by denoising (RED), deep unfolding (DU), compressed sensing using generative models

(CSGM), and deep equilibrium models (DEQ) (Bora et al., 2017; Romano et al., 2017; Zhang & Ghanem, 2018; Hauptmann et al., 2018; Gilton et al., 2021a; Liu et al., 2022; Hu et al., 2024d). These approaches come with different trade-offs in terms of imaging performance, computational and memory complexity, flexibility, need for supervision, and theoretical understanding.

**Denoisers as Priors.** Score-based models (SBMs) are a powerful subset of DL methods for solving inverse problems that use deep Gaussian denoisers as imaging priors. Plug-and-Play (PnP) methods can be viewed as SBMs that incorporate denoisers within iterative optimization algorithms (see recent reviews (Ahmad et al., 2020; Kamilov et al., 2023)). These approaches construct a cost function by combining an explicit likelihood with a score function implicitly defined by the denoiser prior. Over the past few years, numerous variants of PnP have been developed (Venkatakrishnan et al., 2013; Romano et al., 2017; Metzler et al., 2018; Dong et al., 2019; Zhang et al., 2019; Wei et al., 2020; Hurault et al., 2022a), which has motivated an extensive research into their theoretical properties and empirical effectiveness (Chan et al., 2017; Buzzard et al., 2018; Ryu et al., 2019; Sun et al., 2019; Tirer & Giryes, 2019; Teodoro et al., 2019; Sun et al., 2021; Cohen et al., 2021; Fang et al., 2024; Renaud et al., 2024; Hu et al., 2024a; Terris et al., 2024). Diffusion Models (DMs) represent another category of SBMs; they are trained to learn the score function of the underlying probability distribution governed by stochastic differential equations (SDEs) (Ho et al., 2020; Song et al., 2021). Once trained, these models can be used as powerful priors for inverse problems by leveraging their learned score functions. Specifically, pre-trained DMs facilitate posterior sampling by guiding the denoising process to generate data consistent with observed measurements. This approach enables DMs to address inverse problems, often achieving impressive perceptual performance even for highly ill-posed inverse problems (Chung et al., 2023; Zhu et al., 2023; Wang et al., 2023; Feng et al., 2023; Sun et al., 2024; Wu et al., 2024; Song et al., 2024; Hu et al., 2024b; Alçalar & Akçakaya, 2024; Zhao et al., 2024; Rout et al., 2024; Bai et al., 2024; Mardani et al., 2024).

**Restoration Networks as Priors.** In addition to denoiser-based methods, recent work has also considered using restoration models as implicit priors for solving inverse problems (Zhang et al., 2019; Liu et al., 2020; Gilton et al., 2021b; Hu et al., 2024c). It has been observed that pre-trained restoration models can be effective priors for addressing unseen inverse problems, sometimes surpassing traditional denoiser-based approaches (Hu et al., 2024c). However, existing methods present two main limitations. First, existing restoration priors have relied on a fixed prior tailored to a specific degradation. Although deep restoration models can be trained in various settings—such as different

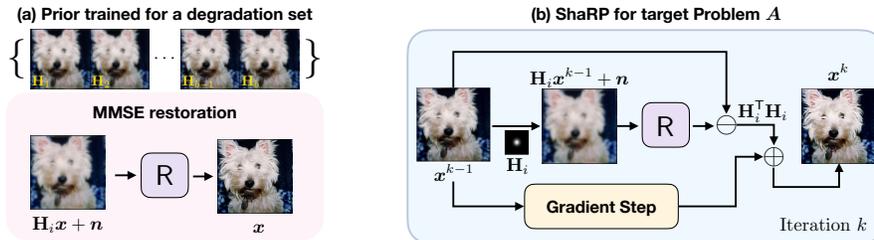


Figure 1: A restoration network trained on a set of tasks  $\{\mathbf{H}_i\}$  can be used as a prior within ShaRP to address different target inverse problems without the need for retraining.

blur kernels for image deblurring or diverse undersampling masks for MRI reconstruction—current approaches do not leverage this capability, limiting their robustness to diverse artifacts. Second, prior work has not explored the potential of learning restoration priors directly from undersampled measurements, without access to fully sampled data. Unlike Gaussian denoisers, training without fully sampled data is natural for restoration models (Ulyanov et al., 2018; Krull et al., 2019; Quan et al., 2020; Yaman et al., 2020; Liu et al., 2020; Tachella et al., 2022; Chen et al., 2022; Millard & Chiew, 2023; Gan et al., 2023b; Terris & Moreau, 2023; Gossard & Weiss, 2024). It is also worth highlighting the related work that has explored using corrupt measurements for training Ambient DMs (Daras et al., 2023; Aali et al., 2024). These Ambient DMs, trained directly on such undersampled measurements, address the reconstruction problem by sampling from their approximated posterior distribution using a standard diffusion sampler at inference. ShaRP, conversely, adopts an optimization approach, defining a novel objective function from an ensemble of likelihoods of multiple degraded observations.

Method	Prior Type	Prior Configuration
PnP	Only Gaussian Denoiser	Single Prior
SNORE	Only Gaussian Denoiser	Ensemble of Priors
DRP	General Restoration <sup>†</sup>	Single Prior
ShaRP	General Restoration <sup>†</sup>	Ensemble of Priors

<sup>†</sup> Can include denoisers and extends to other restoration tasks.

Table 1: ShaRP differs from existing methods by offering greater flexibility through the use of an ensemble of priors trained on general restoration tasks, extending beyond Gaussian denoising priors.

**Our contribution.** (1) We propose ShaRP, a new framework for solving inverse problems leveraging a set of priors implicit in a pre-trained deep restoration network. As summarized in Table 1, ShaRP generalizes Regularization by Denoising (RED) (Romano et al., 2017) and Stochastic Denoising Regularization (SNORE) (Renaud et al., 2024) by using more flexible restoration operators and generalizes Deep Restoration Priors (DRP) (Hu et al., 2024c) by using

### Algorithm 1 Stochastic deep Restoration Priors (ShaRP)

- 1: **input:** Initial value  $x^0 \in \mathbb{R}^n$ ,  $\gamma > 0$ ,  $\sigma > 0$ , and  $\tau > 0$
- 2: **for**  $k = 1, 2, 3, \dots$  **do**
- 3:   Sample a degradation operator:  $\mathbf{H} \sim p(\mathbf{H})$
- 4:    $s \leftarrow \mathbf{H}x^{k-1} + n$  with  $n \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$
- 5:    $x^k \leftarrow x^{k-1} - \gamma \widehat{\nabla} f(x^{k-1})$   
       with  $\widehat{\nabla} f(x) := \nabla g_A(x) + \widehat{\nabla} h(x)$   
       where  $\widehat{\nabla} h(x) := \frac{\tau}{\sigma^2} \mathbf{H}^T \mathbf{H}(x - \mathbf{R}(s, \mathbf{H}))$
- 6: **end for**

multiple restoration priors instead of relying on a single one. (2) We introduce a novel regularization concept for inverse problems that encourages solutions that produce degraded versions closely resembling real degraded images. For example, our regularizer favors an MR image solution only if its various degraded versions are consistent with the characteristics of actual degraded MR images. (3) We show that ShaRP can be interpreted as a stochastic gradient method that seeks to minimize a composite objective that incorporates our proposed regularizer. We discuss its convergence for both exact and approximate MMSE restoration operators. We also provide theoretical guarantees on learning restoration operators from incomplete measurements. (4) We implement ShaRP with both supervised and self-supervised restoration models as priors and test it on two inverse problems: compressed sensing MRI (CS-MRI) and single-image super-resolution (SISR). Our results highlight the capability of restoration models to achieve state-of-the-art performance. Notably, in the MRI context, we show that restoration networks trained directly on subsampled and noisy MRI data can serve as effective priors, a scenario where training traditional Gaussian denoisers is infeasible.

## 3. Stochastic Deep Restoration Priors

ShaRP is presented in Algorithm 1. It considers a prior based on a deep restoration model  $\mathbf{R}(s, \mathbf{H})$  pre-trained using the family of degradation operators, such as blur kernels or MRI masks. More specifically, the deep restoration model  $\mathbf{R}$

is trained to solve the following set of restoration problems

$$\mathbf{s} = \mathbf{H}\mathbf{x} + \mathbf{n} \quad \text{with} \quad \mathbf{x} \sim p_{\mathbf{x}}, \quad \mathbf{H} \sim p(\mathbf{H}), \quad (3)$$

where  $\mathbf{n}$  is the AWGN vector with variance  $\sigma^2$  and  $p_{\mathbf{x}}$  denotes the probability distribution of the target images, and  $p(\mathbf{H})$  is the probability density of considered degradation operators. Importantly, the restoration problems (3) are used exclusively for training R and do not need to match the target inverse problem (1), which involves the measurement operator  $\mathbf{A}$ .

We prove below that  $\widehat{\nabla} f$  corresponds to a stochastic approximation of an objective function of the form  $f = g + h$ . Here,  $g(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2$  and the term  $h(\mathbf{x})$  is defined as  $h(\mathbf{x}) = \tau \mathbb{E}_{\mathbf{s} \sim G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x}), \mathbf{H} \sim p(\mathbf{H})} [-\log p(\mathbf{s}|\mathbf{H})]$ . We will further introduce the derivation and interpretation of  $h(\mathbf{x})$  in Section 4.1. Similar to traditional stochastic gradient methods, ShaRP can be implemented using various selection strategies for the degradation operators.

Each iteration of ShaRP has an intuitive interpretation, where the next solution is obtained by combining the gradient of the data-fidelity term  $\nabla g$  and the residual of restored image corresponding to the selected degradation operator. It is worth highlighting that ShaRP is compatible with multiple priors, each restoring an image from Gaussian or non-Gaussian degradations. Also, the computational cost of running ShaRP is comparable to those of single-model approaches. This is due to the stochastic nature of our algorithm that uses a restoration operator in each iteration.

## 4. Theoretical analysis of ShaRP

We now present several theoretical results on ShaRP. Our main contribution is the expression for the closed-form regularizer minimized by ShaRP. We also show how to interpret ShaRP as a stochastic gradient method, enabling its convergence analysis when using inexact MMSE operators. We finally prove theoretical guarantees on learning MMSE restoration operators from undersampled measurements.

### 4.1. Explicit formula for regularizer

Consider a restoration model that perform MMSE estimation of  $\mathbf{x} \in \mathbb{R}^n$  for problems (3)

$$\begin{aligned} \mathbf{R}^*(\mathbf{s}, \mathbf{H}) &= \mathbb{E}[\mathbf{x} | \mathbf{s}, \mathbf{H}] = \int \mathbf{x} p(\mathbf{x} | \mathbf{s}, \mathbf{H}) d\mathbf{x} \\ &= \frac{1}{p(\mathbf{s} | \mathbf{H})} \int \mathbf{x} G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x}) p_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (4)$$

where we used the probability density of the observation  $\mathbf{s}$  conditioned on the operator  $\mathbf{H}$

$$p(\mathbf{s}|\mathbf{H}) = \int G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x}) p_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}. \quad (5)$$

The function  $G_\sigma$  in (5) denotes the Gaussian density function with the standard deviation  $\sigma > 0$ .

We propose the ShaRP regularizer

$$h(\mathbf{x}) = \tau \mathbb{E}_{\mathbf{s} \sim G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x}), \mathbf{H} \sim p(\mathbf{H})} [-\log p(\mathbf{s}|\mathbf{H})], \quad (6)$$

where  $\tau > 0$  is the regularization parameter and  $p(\mathbf{H})$  is the distribution of considered degradations. The regularizer  $h$  is minimized when degraded versions of  $\mathbf{x}$  are highly probable under the distribution  $p(\mathbf{s}|\mathbf{H})$ , with  $\mathbf{H}$  sampled from  $p(\mathbf{H})$ . In other words, a solution  $\hat{\mathbf{x}}$  is considered effective if its degraded versions  $\mathbf{H}\hat{\mathbf{x}}$  align with the degraded versions  $\mathbf{H}\mathbf{x}$  of clean images  $\mathbf{x} \sim p_{\mathbf{x}}$ , for all  $\mathbf{H} \sim p(\mathbf{H})$ . This can be understood as searching for a fixed point that exhibits equivariance across multiple degradations. The key benefit of the proposed regularizer in (6) lies in its versatility, allowing a wide range of degradation operators to be used in a unified framework. In particular, this formulation remains compatible with Gaussian denoisers, as  $\mathbf{H} = \mathbf{I}$  can always be incorporated into  $p(\mathbf{H})$ .

We are now ready to state our first theoretical result.

**Theorem 1.** *Assume that the prior density  $p_{\mathbf{x}}$  is non-degenerate over  $\mathbb{R}^n$  and let  $\mathbf{R}^*$  be the MMSE restoration operator (4) corresponding to the restoration problems (3). Then, we have that*

$$\nabla h(\mathbf{x}) = \frac{\tau}{\sigma^2} \left( \mathbb{E} \left[ \mathbf{H}^T \mathbf{H} (\mathbf{x} - \mathbf{R}^*(\mathbf{s}, \mathbf{H})) \right] \right), \quad (7)$$

where  $h$  is the ShaRP regularizer in (6), the expectation is with respect to  $\mathbf{s} \sim G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x})$  and  $\mathbf{H} \sim p_{\mathbf{H}}$ .

The proof is in the appendix. Note that the expression within the square parenthesis in (7) matches the ShaRP update in Line 4 of Algorithm 1, which directly implies that ShaRP using the exact MMSE restoration operator  $\mathbf{R}^*$  is a stochastic gradient method for minimizing  $f = g + h$ , where  $g$  is the data-fidelity term and  $h$  is the ShaRP regularizer in (6).

### 4.2. Theoretical convergence of ShaRP

Given the explicit regularized in (6), the iterations of ShaRP in Algorithm 1 can be seen as stochastic gradient method for minimizing  $f = g + h$ . In practical scenarios, the learned restoration model may be imperfect, meaning it cannot be considered a perfect MMSE estimator. We now present the convergence analysis of ShaRP under a restoration operator  $\mathbf{R}$  that *approximates* the true MMSE restoration operator  $\mathbf{R}^*$ . We adopt the analysis of biased stochastic gradient descent (SGD) from the optimization literature (Bertsekas, 2011; Ghadimi & Lan, 2016; Demidovich et al., 2023), by interpreting iterations of ShaRP as a variant of biased SGD. For a given degraded observation  $\mathbf{s} = \mathbf{H}\mathbf{x} + \mathbf{n}$  with

$\mathbf{H} \sim p_{\mathbf{H}}$  and  $\mathbf{n} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ , we define the stochastic gradient used by ShaRP

$$\begin{aligned} \widehat{\nabla} f(\mathbf{x}) &= \nabla g(\mathbf{x}) + \widehat{\nabla} h(\mathbf{x}), \\ \text{where } \widehat{\nabla} h(\mathbf{x}) &:= \frac{\tau}{\sigma^2} \mathbf{H}^T \mathbf{H} (\mathbf{x} - \mathbf{R}(\mathbf{s}, \mathbf{H})). \end{aligned} \quad (8)$$

Since  $\mathbf{R}$  is an inexact MMSE restoration operator, we also define the bias vector

$$\mathbf{b}(\mathbf{x}) = \frac{\tau}{\sigma^2} \mathbb{E} \left[ \mathbf{H}^T \mathbf{H} (\mathbf{R}^*(\mathbf{s}, \mathbf{H}) - \mathbf{R}(\mathbf{s}, \mathbf{H})) \right] \quad (9)$$

which quantifies the average difference between the exact and inexact MMSE restoration operators with respect to  $\mathbf{H} \sim p_{\mathbf{H}}$  and  $\mathbf{s} \sim G_{\sigma}(\mathbf{s} - \mathbf{H}\mathbf{x})$ . The analysis requires three assumptions that jointly serve as sufficient conditions for our theorem.

**Assumption 1.** *The function  $f$  has a finite minimum  $f^* > -\infty$  and the gradient  $\nabla f$  is Lipschitz continuous with constant  $L > 0$ .*

This is a standard assumption used in the analysis of gradient-based algorithms (see (Nesterov, 2004), for example). It is satisfied by a large number of functions, including the traditional least-squares data-fidelity function.

**Assumption 2.** *The stochastic gradient has a bounded variance for all  $\mathbf{x} \in \mathbb{R}^n$ , which means that there exists a constant  $\nu > 0$  such that*

$$\mathbb{E} \left[ \left\| \widehat{\nabla} f(\mathbf{x}) - \mathbb{E} \left[ \widehat{\nabla} f(\mathbf{x}) \right] \right\|_2^2 \right] \leq \nu^2,$$

where expectations are with respect to  $\mathbf{H} \sim p_{\mathbf{H}}$  and  $\mathbf{s} \sim G_{\sigma}(\mathbf{s} - \mathbf{H}\mathbf{x})$ .

This is another standard assumption extensively used in the analysis of online or stochastic optimization algorithms (Bertsekas, 2011; Ghadimi & Lan, 2016; Demidovich et al., 2023).

**Assumption 3.** *The bias  $\mathbf{b}(\mathbf{x})$ , as defined in (9), is bounded, which means that there exists  $\varepsilon > 0$  such that for all  $\mathbf{x} \in \mathbb{R}^n$*

$$\|\mathbf{b}(\mathbf{x})\|_2 \leq \varepsilon.$$

The only assumption on the bias is that it is bounded, which is a relatively mild assumption, as image pixel values are typically constrained (e.g., to  $[0, 255]$ ).

Note that Assumptions 1-2 are needed only for Proposition 1. They are not needed for our main results—Theorem 1-2.

**Proposition 1.** *Run ShaRP for  $t \geq 1$  iterations using the step-size  $0 < \gamma \leq 1/L$  under Assumptions 1-3. Then, the sequence  $\mathbf{x}^k$  generated by ShaRP satisfies*

$$\mathbb{E} \left[ \frac{1}{t} \sum_{k=1}^t \|\nabla f(\mathbf{x}^{k-1})\|_2^2 \right] \leq \frac{2}{\gamma t} (f(\mathbf{x}^0) - f^*) + \gamma L \nu^2 + \varepsilon^2.$$

The general analysis of biased stochastic SGD has been extensively discussed in the optimization literature (Bertsekas, 2011; Ghadimi & Lan, 2016; Demidovich et al., 2023). Our contribution here is the relationship between the iterates of ShaRP and those of biased SGD, leading to theoretical guarantees on the stability of ShaRP using restoration networks that do not correspond to ideal MMSE estimators. Proposition 1 states that *in expectation*, ShaRP minimizes the norm of the gradient  $\nabla f$  up to an error term that has two components,  $\gamma L \nu^2$  and  $\varepsilon^2$ . Since the first component depends on  $\gamma$ , it can be made as small as desired by controlling the step-size  $\gamma$ . The second component only depends on the magnitude of the bias  $\varepsilon$ , which, in turn, directly depends on the accuracy of the restoration operator relative to the true MMSE restoration operator  $\mathbf{R}^*$ .

### 4.3. Learning restoration priors without groundtruth

In this subsection, we present a theoretical result that establishes the feasibility of learning an MMSE estimator from undersampled measurements. Let the undersampled measurements be defined as  $\mathbf{s} = \mathbf{H}\mathbf{x} + \mathbf{n}$ , where  $\mathbf{H} = \mathbf{P}\mathbf{M}$ . Here,  $\mathbf{P}$  represents a binary subsampling matrix,  $\mathbf{M}$  denotes the square transfer operator, and  $\mathbf{n}$  corresponds to the noise vector.

To show that an MMSE estimator can be learned from undersampled measurement, we need the following assumption.

**Assumption 4.**  $\mathbb{E}_{\mathcal{P}}[\mathbf{P}^T \mathbf{P}]$  has a full rank where the expectation is taken over  $p_{\mathcal{P}}$ . The measurement operator  $\mathbf{M}$  is orthogonal matrix.

This assumption implies that the *union* of all sampling matrices  $\mathbf{P}$  spans the complete measurement domain, even though each individual  $\mathbf{P}$  may remain undersampled.

**Theorem 2.** *Under Assumption 4, the MMSE estimator  $\mathbf{R}$  learned using the weighted self-supervised loss ( $\ell_{\text{self}}$ ) is equivalent to its supervised counterpart ( $\ell_{\text{sup}}$ ). Specifically, we have:*

$$\mathbf{R}_{\ell_{\text{self}}}(\boldsymbol{\theta}) = \mathbf{R}_{\ell_{\text{sup}}}(\boldsymbol{\theta}). \quad (10)$$

where

$$\ell_{\text{sup}} = \mathbb{E} \left[ \frac{1}{2} \|\bar{\mathbf{x}} - \mathbf{x}\|_2^2 \right] \quad (11)$$

and

$$\ell_{\text{self}} = \mathbb{E} \left[ \frac{1}{2} \|\mathbf{P}' \mathbf{M} \bar{\mathbf{x}} - \mathbf{s}'\|_{\mathbf{W}}^2 \right], \quad (12)$$

where the vector  $\bar{\mathbf{x}} = \mathbf{R}(\mathbf{s})$  is MMSE estimation of  $\mathbf{R}$  for  $\mathbf{s}$ ,  $\mathbf{s}' = \mathbf{P}' \mathbf{M} \mathbf{x}$  is another independently subsampled measurement and  $\mathbf{W}$  is a weighting matrix to compensate sampling imbalance in measurement domain.

## 5. Numerical Results

We numerically validate ShaRP on two inverse problems of the form  $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$ : (Compressive Sensing MRI (CS-MRI) and (b) Single Image Super Resolution (SISR). In both cases,  $\mathbf{e}$  represents additive white Gaussian noise (AWGN). For the data-fidelity term in eq. (2), we use the  $\ell_2$ -norm loss for both problems. Quantitative performance is evaluated by Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM). Additionally, for the SISR task, we include the Learned Perceptual Image Patch Similarity (LPIPS) metric to evaluate perceptual quality. Additional numerical results are provided in the supplementary material.

Noise level	$\sigma = 0.005$		$\sigma = 0.010$		$\sigma = 0.015$	
Metrics	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Zero-filled	26.93	0.848	26.92	0.847	26.90	0.848
TV	31.17	0.923	31.08	0.921	30.91	0.915
PnP-FISTA	35.88	0.938	31.14	0.894	30.32	0.846
PnP-ADMM	<u>35.76</u>	<u>0.941</u>	32.36	0.878	30.66	0.838
DRP	35.52	0.936	32.32	0.914	30.57	0.901
DPS	32.62	0.888	31.39	0.870	30.29	0.856
DDS	35.21	0.937	<u>35.03</u>	<u>0.935</u>	<u>34.51</u>	<u>0.925</u>
ShaRP	<b>37.59</b>	<b>0.963</b>	<b>35.81</b>	<b>0.951</b>	<b>34.92</b>	<b>0.942</b>
E2E-VarNet <sup>†</sup>	38.10	0.971	36.80	0.967	35.79	0.954

<sup>†</sup>E2E-VarNet needs to be retrained for each task.

Table 2: Quantitative comparison of ShaRP with several baselines for CS-MRI using  $4\times$  uniform masks. The **best** and **second best** results across general image restoration methods are highlighted. Notably, ShaRP outperforms SOTA baseline methods.

Noise level	$\sigma = 0.005$		$\sigma = 0.010$		$\sigma = 0.015$	
Metrics	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
ADMM-TV	28.14	0.866	28.06	0.863	27.96	0.859
GRAPPA	28.09	0.792	25.39	0.699	23.94	0.649
SPICER	<u>31.87</u>	<u>0.901</u>	<u>31.67</u>	<u>0.889</u>	<u>31.50</u>	<u>0.887</u>
ShaRP <sup>self</sup>	<b>33.87</b>	<b>0.909</b>	<b>33.64</b>	<b>0.900</b>	<b>33.21</b>	<b>0.892</b>

Table 3: Quantitative comparison of ShaRP with a self-supervised pre-trained restoration operator, compared to several baselines for CS-MRI using  $4\times$  uniform masks. The **best** and **second best** results are highlighted. Note the excellent performance of ShaRP even using priors trained without fully-sampled ground-truth data.

### 5.1. CS-MRI setting

The complex-valued measurement model for CS-PMRI is expressed as:  $\mathbf{y} = \mathbf{PFS}\mathbf{x} + \mathbf{e}$ . Here,  $\mathbf{x}$  represents the underlying complex-valued image,  $\mathbf{S}$  are the complex multi-coil sensitivity maps,  $\mathbf{F}$  is the Fourier transform operator,

and  $\mathbf{P}$  is the k-space subsampling mask. The term  $\mathbf{y}$  denotes the acquired complex k-space data, and  $\mathbf{e}$  is the additive complex noise vector. We utilized the open-access fastMRI dataset; further experimental details can be found in Section B.1 of the supplementary material.

**Ensemble of restoration priors for CS-MRI.** Recent methods like InDI (Delbracio & Milanfar, 2023) and I<sup>2</sup>SB (Liu et al., 2023) use controllable processes to train ensembles of restoration priors, each an MMSE operator for a specific setting. We build on this by training an  $8\times$  uniform subsampling CS-MRI model with 8 masks as our prior. Following InDI, we decompose the MRI degradation operator  $\mathbf{M} = \mathbf{PFS}$  into convex combinations of  $\mathbf{M}$  and the identity  $\mathbf{I}$ :  $\mathbf{H}_\alpha = (1 - \alpha)\mathbf{I} + \alpha\mathbf{M}$ , with  $\alpha$  controlling degradation. Varying  $\alpha$  creates an ensemble of tasks. Training R on all these tasks allows it to act as an ensemble of MMSE operators:  $R(\mathbf{s}, \mathbf{H}_\alpha) = \mathbb{E}[\mathbf{x}|\mathbf{s}, \mathbf{H}_\alpha]$ . We used the MSE loss for training.

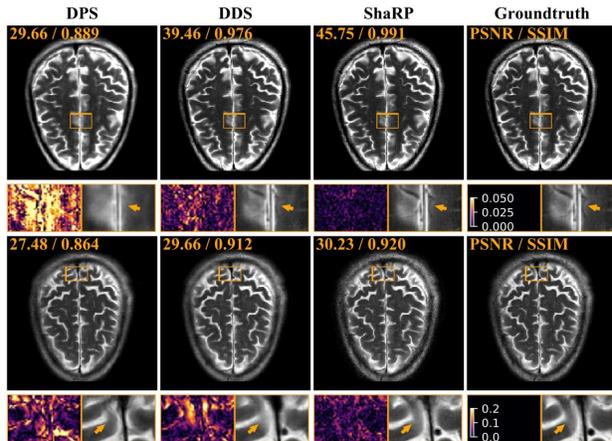


Figure 2: Visual comparison of ShaRP with baseline methods on CS-MRI. The top row shows results for a  $4\times$  random mask with noise  $\sigma = 0.005$ , and the bottom row for a  $6\times$  random mask with noise  $\sigma = 0.015$ . PSNR and SSIM values are in the top-left corner of each image. Error maps and zoomed-in areas highlight differences. Notably, ShaRP with stochastic priors outperforms state-of-the-art methods using denoiser and diffusion model priors.

**Training restoration priors without groundtruth.** When fully-sampled images are unavailable for training, MRI restoration priors can be trained self-supervisedly (Yaman et al., 2020; Millard & Chiew, 2023; Gan et al., 2023b; Hu et al., 2024d), as demonstrated in Theorem 2. Notably, the MRI imaging system aligns well with the assumptions required by Theorem 2, as  $\mathbf{FS}$  represents an orthogonal matrix. This approach uses a separate subsampled measurement as the training label, instead of the ground-truth image. Importantly, and consistent with Theorem 2, our self-supervised training *theoretically* achieves MMSE restora-

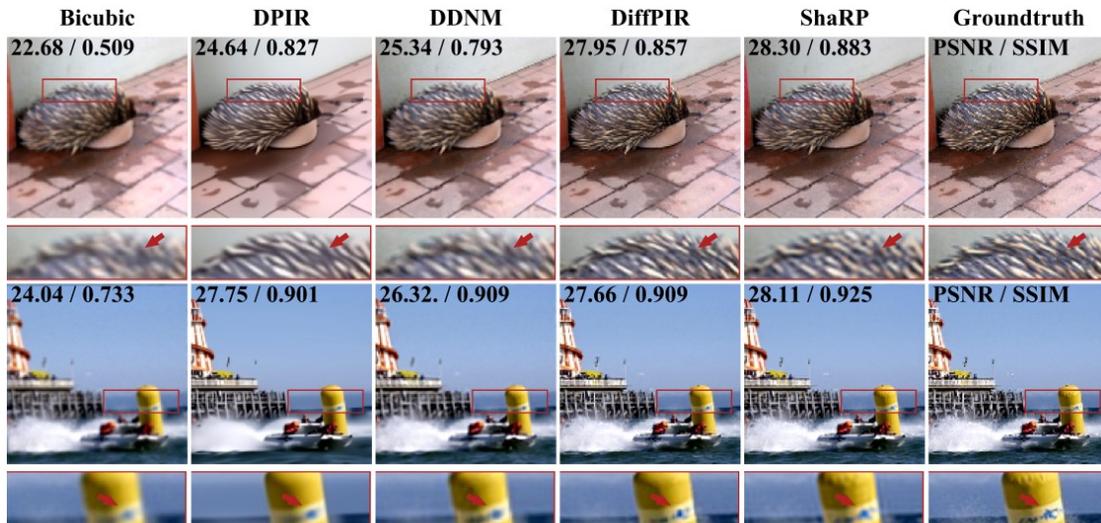


Figure 3: Visual comparison of ShaRP with several well-known methods on  $2\times$  SISR. The top row shows results for SISR with gaussian blur kernel with  $\sigma = 1.25$ , while the bottom row shows results for SISR with gaussian blur kernel with  $\sigma = 1.5$ . The quantities in the top-left corner of each image provide PSNR and SSIM values for each method. The squares at the bottom of each image visualize the zoomed area in the image.

tion priors equivalent to those trained fully-supervisedly. For validation, we train an  $8\times$  uniform subsampling CS-MRI model to handle eight distinct restoration operators corresponding to different sampling masks. Additional training details are in Section B.1 of the supplement.

With the pre-trained  $8\times$  models as ensembles of restoration priors, we evaluate ShaRP’s performance across a variety of configurations, including two sub-sampling rates ( $4\times$  and  $6\times$ ), two mask types (uniform and random), and three noise levels ( $\sigma = 0.005, 0.01, \text{ and } 0.015$ ). Due to space constraints, only a subset of these settings is included in the main paper. Additional results can be found in Section B.1 of the supplementary material.

**Baselines.** ShaRP was compared against several baseline methods, including denoiser-based approaches (PnP-FISTA (Kamilov et al., 2017), PnP-ADMM (Chan et al., 2017)) and diffusion model-based methods (DPS (Chung et al., 2023), DDS (Chung et al., 2024)). To highlight the advantages of using a stochastic set of restoration operators, we also compared ShaRP with the DRP method (Hu et al., 2024c), which applies only a single restoration operator. Additional details related to the baseline methods can be found in Section B.1 of the appendix.

**Results with supervised MMSE restoration operator.** Table 2 provides a quantitative comparison of reconstruction performance across different acceleration factors and noise levels using a uniform sub-sampling mask. In all configurations, ShaRP consistently outperforms the baseline methods. The use of a set of restoration operators clearly enhances

ShaRP’s performance, highlighting the effectiveness of employing multiple operators to maximize the regularization information provided by the restoration model. Figure 2 presents visual reconstructions for two test scenarios, where ShaRP accurately recovers fine brain details, particularly in the zoomed-in regions, while baseline methods tend to oversmooth or introduce artifacts. These results highlight ShaRP’s superior ability to manage structured artifacts and preserve fine details, outperforming both denoiser-based and diffusion model-based methods.

**Results with MMSE restoration operator learned from incomplete measurements.** We further evaluate ShaRP’s performance using an restoration model, learned in a self-supervised manner, as introduced in (Gan et al., 2023b). In this setting, we compare ShaRP against two classical methods for CS-MRI reconstruction without groundtruth: TV (Block et al., 2007) and GRAPPA (Griswold et al., 2002) and a recent state-of-the-art self-supervised deep unrolling method: SPICER (Hu et al., 2024d). As shown in Table 3, ShaRP demonstrates its effectiveness in leveraging a self-supervised restoration prior for various reconstruction tasks, even when only incomplete measurements ( $8\times$  subsampled) are available. Note that, ShaRP using self-supervised restoration prior even outperforms DPIP and DPS that use Gaussian denoisers trained using fully-sampled ground truth images (see Table 6 in the appendix). It is important to highlight that training Gaussian denoisers is infeasible when only undersampled measurements are available.

We further evaluate the convergence performance of ShaRP using both supervised and self-supervised restoration priors.

Due to space limitations, the detailed results and discussion are provided in Section C.3.

## 5.2. Single Image Super Resolution (SISR)

The measurement operator in SISR can be written as  $A = SK$ , where  $K$  represents convolution with the blur kernel, and  $S$  performs standard  $d$ -fold down-sampling. In our experiments, we use two Gaussian blur kernels  $k$ , each with distinct standard deviations (1.25 and 1.5), and with down-sampling factor of 2. Both noisy and noise-free cases are considered to evaluate the noisy robustness of ShaRP. We randomly selected 100 images from the ImageNet test set, as provided in DiffPIR<sup>1</sup>. Due to space constraints, only a subset of comparisons is included in the main paper. Additional results can be found in Section D of the supplement.

### Ensemble of Restoration Priors for Image Deblurring.

Similar to our CS-MRI prior training, we decompose the deblurring task (Gaussian blur operator  $K$ ) into a convex combination of  $K$  and the identity mapping  $I$ :  $H_\alpha = (1 - \alpha)I + \alpha K$ , with  $\alpha$  controlling degradation. Varying  $\alpha$  yields multiple degradation operators, enabling our restoration network  $R$  to handle each. This allows  $R$  to function as an ensemble of MMSE restoration operators:  $R(s, H_\alpha) = \mathbb{E}[x | s, H_\alpha]$ , where  $s$  and  $x$  are the degraded and original images, respectively. The original  $K$  uses a  $31 \times 31$  Gaussian kernel with a standard deviation of 3. Further training details are in Section B.2 of the supplement.

**Baselines.** We compared ShaRP against baselines including DPIR (Zhang et al., 2022), a state-of-the-art PnP method using pre-trained denoisers, and diffusion-based methods DPS (Chung et al., 2023), DDNM (Wang et al., 2023), DDRM (Kawar et al., 2022), and DiffPIR (Zhu et al., 2023), which use different sampling strategies for SISR.

**Results on SISR with deblurring prior.** Figure 3 shows the visual reconstruction results for two settings with different blur kernels. As demonstrated, ShaRP successfully recovers most features and maintains high data consistency with the available measurements. Table 4 provides quantitative comparisons against baselines across blur kernels and noise levels. ShaRP achieves the highest PSNR and SSIM, but ranks second in perceptual performance (LPIPS), consistent with SOTA Diffusion Model-based methods on SISR. However, the deblurring prior in ShaRP recovers fine details, ensuring competitive perceptual quality. One interpretation for the improved performance is that when considering an ensemble of blur-then-deblur operations, ShaRP can be seen as masking and subsequently restoring portions of null-space—akin to the mechanism of masked (denoising) autoencoders. In this interpretation, MMSE restoration networks effectively reconstruct the missing null-space signals,

which improves the perceptual quality. Additional results can be found in Section D of supplementary material.

## 6. Discussion

**ShaRP as a Unified Framework.** Unlike methods that rely on a fixed prior for all tasks, ShaRP operates without being constrained to any specific prior. Instead, it adapts flexibly to the unique characteristics of diverse inverse problems, enabling the integration of a broad range of restoration priors, including—but not limited to—denoisers. This adaptability not only enhances the versatility and applicability of the framework but also has the potential to achieve superior performance, as demonstrated in Section 5. By building on a unified theoretical foundation, ShaRP represents a significant advancement in addressing a wide variety of inverse problems effectively.

**Performance improved with Ensemble Priors.** Conventional ensemble priors for inverse problems typically focus on Gaussian denoisers across varying noise levels (Zhang et al., 2019; Hurault et al., 2022a; Renaud et al., 2024). ShaRP advances this concept by integrating structured noise from diverse degradation operators, forming a richer and more robust ensemble. Empowered by its flexible and novel regularizer, ShaRP seamlessly incorporates this diverse ensemble into a unified, closed-form prior that accommodates both Gaussian and more complex denoisers. This innovation highlights ShaRP’s ability to leverage a broader range of priors, enabling adaptable and robust performance across a wide variety of noise environments. Further evidence of the benefits of such an expansive ensemble is provided by an ablation study in Section E.2.

ShaRP’s ability to integrate diverse restoration priors within a unified framework enables task-specific adaptability, driving significant performance improvements and advancing inverse problem-solving.

**Sampling from Restoration Priors.** Our current work introduces a novel objective function balancing data fidelity with implicit regularization from an ensemble of priors. Although this optimization-based approach achieves high reconstruction accuracy, its perceptual performance is inherently limited by the well-known perceptual-distortion trade-off (Blau & Michaeli, 2018). To enhance the perceptual performance, sampling-based extension is a promising future direction. Specifically, we aim to adapt existing diffusion-based image restoration methods (Daras et al., 2023; Bansal et al., 2023; Delbraccio & Milanfar, 2023; Liu et al., 2023), initially developed for specific tasks, to serve as priors for sampling in general image restoration.

<sup>1</sup><https://github.com/yuanzhi-zhu/DiffPIR/tree/main/testsets>



Noise level	Noiseless			$\sigma = 0.01$			Noiseless			$\sigma = 0.01$		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
DPIR	28.10	0.809	0.305	28.05	0.807	0.308	27.90	0.803	0.314	27.87	0.800	0.314
DDNM	27.53	0.786	0.240	27.49	0.784	0.246	27.02	0.764	0.264	27.01	0.763	0.267
DPS	24.68	0.661	0.395	24.60	0.657	0.399	24.50	0.657	0.403	24.44	0.655	0.406
DiffPIR	28.92	0.852	<b>0.152</b>	28.63	0.839	<b>0.169</b>	<u>28.59</u>	<u>0.834</u>	<b>0.172</b>	<u>28.02</u>	<u>0.819</u>	<b>0.185</b>
DDRM	28.20	0.845	0.161	28.11	0.832	0.188	27.93	0.826	0.188	27.67	0.817	0.193
DRP	<u>29.28</u>	<u>0.868</u>	0.207	<u>28.87</u>	<u>0.848</u>	0.248	28.24	0.836	0.235	28.01	0.820	0.278
ShaRP	<b>30.09</b>	<b>0.891</b>	<u>0.179</u>	<b>29.03</b>	<b>0.852</b>	<u>0.223</u>	<b>29.28</b>	<b>0.872</b>	<u>0.209</u>	<b>28.06</b>	<b>0.821</b>	<u>0.268</u>

Table 4: Quantitative comparison of ShaRP with several baselines for SISR based on two different blur kernels on ImageNet dataset. The **best** and **second best** results are highlighted. Notably, ShaRP outperforms SOTA methods based on denoisers and diffusion models.

### 7. Conclusion

This work presents ShaRP, a novel framework for imaging inverse problems that leverages pre-trained restoration networks as priors. Unlike methods limited to Gaussian denoisers or a single prior, ShaRP integrates multiple priors tailored to diverse degradations within a unified theoretical framework. The key findings are: (1) using priors beyond traditional Gaussian denoisers broadens the framework’s applicability and performance, and (2) stochastically integrating multiple degradation-specific priors achieves better performance than relying on a single prior. Numerical experiments confirm ShaRP’s superiority over conventional methods. By unifying diverse priors into a flexible framework, ShaRP encourages exploring more complex restoration priors for inverse problem-solving.

### Ethics Statement

To the best of our knowledge this work does not give rise to any significant ethical concerns.

### Impact Statement

ShaRP makes advanced computational imaging more accessible by reducing reliance on task-specific models and extensive supervision. Its versatility enables broad societal impact across fields that depend on high-quality image data, such as healthcare, scientific research, and consumer technology.

### Acknowledgement

Research presented in this article was supported in part by the NSF CAREER awards under grant CCF-2043134.

### References

Aali, A., Daras, G., Levac, B., Kumar, S., Dimakis, A. G., and Tamir, J. I. Ambient diffusion posterior sampling: Solving inverse problems with diffusion models trained on corrupted data. *arXiv preprint arXiv:2403.08728*, 2024.

Ahmad, R., Bouman, C. A., Buzzard, G. T., Chan, S., Liu, S., Reehorst, E. T., and Schniter, P. Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery. *IEEE Sig. Process. Mag.*, 37(1):105–116, 2020.

Alçalar, Y. U. and Akçakaya, M. Zero-shot adaptation for approximate posterior sampling of diffusion models in inverse problems. *arXiv preprint arXiv:2407.11288*, 2024.

Bai, W., Chen, S., Chen, W., and Sun, H. Blind inversion using latent diffusion priors. *arXiv preprint arXiv:2407.01027*, 2024.

Bansal, A., Borgnia, E., Chu, H., Li, J., Kazemi, H., Huang, F., Goldblum, M., Geiping, J., and Goldstein, T. Cold diffusion: Inverting arbitrary image transforms without noise. *Proc. Adv. Neural Inf. Process. Syst.*, 36:41259–41282, 2023.

Bertsekas, D. P. Incremental proximal methods for large scale convex optimization. *Math. Program. Ser. B*, 129: 163–195, 2011.

Blau, Y. and Michaeli, T. The perception-distortion tradeoff. In *Proc. IEEE Conf. Comput. Vis. and Pattern Recognit. (CVPR)*, pp. 6228–6237, 2018.

Block, K. T., Uecker, M., and Frahm, J. Undersampled radial MRI with multiple coils. iterative image reconstruction using a total variation constraint. *Magn. Reson. Med.*, 2007.

- Bora, A., Jalal, A., Price, E., and Dimakis, A. G. Compressed sensing using generative priors. In *Int. Conf. Mach. Learn.*, pp. 537–546, August 2017.
- Buzzard, G. T., Chan, S. H., Sreehari, S., and Bouman, C. A. Plug-and-play unplugged: Optimization free reconstruction using consensus equilibrium. *SIAM J. Imaging Sci.*, 11(3):2001–2020, Sep. 2018.
- Chan, S. H., Wang, X., and Elgendy, O. A. Plug-and-play ADMM for image restoration: Fixed-point convergence and applications. *IEEE Trans. Comp. Imag.*, 3(1):84–98, Mar. 2017.
- Chen, D., Tachella, J., and Davies, M. Robust equivariant imaging: a fully unsupervised framework for learning to image from noisy and partial measurements. In *Proc. IEEE Conf. Comput. Vis. and Pattern Recognit. (CVPR)*, 2022.
- Chen, H., Zhang, Y., Kalra, M. K., Lin, F., Chen, Y., Liao, P., Zhou, J., and Wang, G. Low-dose CT with a residual encoder-decoder convolutional neural network. *IEEE Trans. Med. Imag.*, 36(12):2524–2535, Dec. 2017.
- Chung, H., Kim, J., Mccann, M. T., K., M. L., and Ye, J. C. Diffusion posterior sampling for general noisy inverse problems. In *Int. Conf. on Learn. Represent.*, 2023.
- Chung, H., Lee, S., and Ye, J. C. Decomposed diffusion sampler for accelerating large-scale inverse problems. In *Proc. Int. Conf. on Learn. Represent. (ICLR)*, 2024.
- Cohen, R., Blau, Y., Freedman, D., and Rivlin, E. It has potential: Gradient-driven denoisers for convergent solutions to inverse problems. In *Proc. Adv. Neural Inf. Process. Syst.* 34, 2021.
- Daras, G., Shah, K., Dagan, Y., Gollakota, A., Dimakis, A., and Klivans, A. Ambient diffusion: Learning clean distributions from corrupted data. In *Proc. Adv. Neural Inf. Process. Syst.*, 2023.
- Daras, G., Chung, H., Lai, C., Mitsufuji, Y., Ye, J. C., Milanfar, P., Dimakis, A. G., and Delbracio, M. A survey on diffusion models for inverse problems. *arXiv preprint arXiv:2410.00083*, 2024.
- Delbracio, M. and Milanfar, P. Inversion by direct iteration: An alternative to denoising diffusion for image restoration. *Trans. on Mach. Learn. Research*, 2023. ISSN 2835-8856.
- Delbracio, M., Talebei, H., and Milanfar, P. Projected distribution loss for image enhancement. In *2021 Int. Conf. on Comput. Photography (ICCP)*, pp. 1–12, 2021.
- Demidovich, Y., Malinovsky, G., Sokolov, I., and Richtarik, P. A guide through the zoo of biased SGD. In *Proc. Adv. Neural Inf. Process. Syst.*, 2023.
- Dong, W., Wang, P., Yin, W., Shi, G., Wu, F., and Lu, X. Denoising prior driven deep neural network for image restoration. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(10):2305–2318, Oct 2019.
- Fang, Z., Buchanan, S., and Sulam, J. What’s in a prior? Learned proximal networks for inverse problems. *Proc. Int. Conf. on Learn. Represent. (ICLR)*, 2024.
- Feng, B. T., Smith, J., Rubinstein, M., Chang, H., Bouman, K. L., and Freeman, W. Score-based diffusion models as principled priors for inverse imaging. In *Proc. IEEE Int. Conf. Comp. Vis. (ICCV)*, pp. 10520–10531, 2023.
- Gan, W., Shoushtari, S., Hu, Y., Liu, J., An, H., and Kamilov, U. S. Block coordinate plug-and-play methods for blind inverse problems. In *Proc. Adv. Neural Inf. Process. Syst.*, 2023a.
- Gan, W., Ying, C., Boroojeni, P. E., Wang, T., Eldeniz, C., Hu, Y., Liu, J., Chen, Y., An, H., and Kamilov, U. S. Self-supervised deep equilibrium models with theoretical guarantees and applications to mri reconstruction. *IEEE Trans. Comp. Imag.*, 2023b.
- Ghadimi, S. and Lan, G. Accelerated gradient methods for nonconvex nonlinear and stochastic programming. *Math. Program. Ser. A*, 156(1):59–99, Mar. 2016.
- Gilton, D., Ongie, G., and Willett, R. Deep equilibrium architectures for inverse problems in imaging. *IEEE Trans. Comput. Imag.*, 7:1123–1133, 2021a.
- Gilton, D., Ongie, G., and Willett, R. Model adaptation for inverse problems in imaging. *IEEE Trans. Comput. Imag.*, 7:661–674, July 2021b.
- Gossard, A. and Weiss, P. Training adaptive reconstruction networks for blind inverse problems. *SIAM J. Imaging Sci.*, 17(2):1314–1346, 2024.
- Griswold, M. A., Jakob, P. M., Heidemann, R. M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., and Haase, A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn. Reson. in Med.*, 47(6):1202–1210, 2002.
- Hauptmann, A., Lucka, F., Betcke, M., Huynh, N., Adler, J., Cox, B., Beard, P., Ourselin, S., and Arridge, S. Model-based learning for accelerated, limited-view 3-D photoacoustic tomography. *IEEE Trans. Med. Imag.*, 37(6):1382–1393, 2018.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Proc. Adv. Neural Inf. Process. Syst.*, 33:6840–6851, 2020.

- Hu, J., Gan, W., Sun, Z., An, H., and Kamilov, U. A plug-and-play image registration network. In *Proc. Int. Conf. on Learn. Represent. (ICLR)*, 2024a.
- Hu, J., Song, B., Xu, X., Shen, L., and Fessler, J. A. Learning image priors through patch-based diffusion models for solving inverse problems. *arXiv preprint arXiv:2406.02462*, 2024b.
- Hu, Y., Delbracio, M., Milanfar, P., and Kamilov, U. A Restoration Network as an Implicit prior. In *Proc. Int. Conf. on Learn. Represent. (ICLR)*, 2024c.
- Hu, Y., Gan, W., Ying, C., Wang, T., Eldeniz, C., Liu, J., Chen, Y., An, H., and Kamilov, U. S. Spicer: Self-supervised learning for mri with automatic coil sensitivity estimation and reconstruction. *Magn. Reson. Med.*, 2024d.
- Hurault, S., Leclaire, A., and Papadakis, N. Gradient step denoiser for convergent plug-and-play. In *Proc. Int. Conf. on Learn. Represent. (ICLR)*, Kigali, Rwanda, May 1-5, 2022a.
- Hurault, S., Leclaire, A., and Papadakis, N. Proximal denoiser for convergent plug-and-play optimization with nonconvex regularization. In *Int. Conf. Mach. Learn.*, pp. 9483–9505, Kigali, Rwanda, 2022b. PMLR.
- Jin, K. H., McCann, M. T., Froustey, E., and Unser, M. Deep convolutional neural network for inverse problems in imaging. *IEEE Trans. Image Process.*, 26(9):4509–4522, Sep. 2017.
- Kadkhodaie, Z. and Simoncelli, E. P. Stochastic solutions for linear inverse problems using the prior implicit in a denoiser. In *Proc. Adv. Neural Inf. Process. Syst.* 34, pp. 13242–13254, December 6-14, 2021.
- Kamilov, U. S., Mansour, H., and Wohlberg, B. A plug-and-play priors approach for solving nonlinear imaging inverse problems. *IEEE Signal. Proc. Let.*, 24(12):1872–1876, Dec. 2017.
- Kamilov, U. S., Bouman, C. A., Buzzard, G. T., and Wohlberg, B. Plug-and-play methods for integrating physical and learned models in computational imaging. *IEEE Signal Process. Mag.*, 40(1):85–97, January 2023.
- Kang, E., Min, J., and Ye, J. C. A deep convolutional neural network using directional wavelets for low-dose x-ray CT reconstruction. *Med. Phys.*, 44(10):e360–e375, 2017.
- Kawar, B., Elad, M., Ermon, S., and Song, J. Denoising diffusion restoration models. *Proc. Adv. Neural Inf. Process. Syst.*, 35:23593–23606, 2022.
- Krull, A., Buchholz, T., and Jug, F. Noise2void-learning denoising from single noisy images. In *Proc. IEEE Conf. Comput. Vis. and Pattern Recognit. (CVPR)*, pp. 2129–2137, 2019.
- Laumont, R., De Bortoli, V., Almansa, A., Delon, J., Durmus, A., and Pereyra, M. Bayesian imaging using plug & play priors: When Langevin meets Tweedie. *SIAM J. Imaging Sci.*, 15(2):701–737, 2022.
- Liang, J., Cao, J., Sun, G., Zhang, K., Van G., L., and Timofte, R. Swinir: Image restoration using swin transformer. In *Proc. IEEE Int. Conf. Comp. Vis. (ICCV)*, pp. 1833–1844, 2021.
- Liu, G., Vahdat, A., Huang, D., Theodorou, E. A., Nie, W., and Anandkumar, A. I2sb: image-to-image schrödinger bridge. In *Proc. Int. Conf. Machine Learning (ICML)*, pp. 22042–22062, 2023.
- Liu, J., Sun, Y., Eldeniz, C., Gan, W., An, H., and Kamilov, U. S. RARE: Image reconstruction using deep priors learned without ground truth. *IEEE J. Sel. Topics Signal Process.*, 14(6):1088–1099, Oct. 2020.
- Liu, J., Xu, X., Gan, W., Shoushtari, S., and Kamilov, U. S. Online deep equilibrium learning for regularization by denoising. In *Proc. Adv. Neural Inf. Process. Syst.*, 2022.
- Mardani, M., Song, J., K., J., and Vahdat., A. A variational perspective on solving inverse problems with diffusion models. In *Proc. Int. Conf. on Learn. Represent. (ICLR)*, 2024.
- McCann, M. T., Jin, K. H., and Unser, M. Convolutional neural networks for inverse problems in imaging: A review. *IEEE Signal Process. Mag.*, 34(6):85–95, 2017.
- Metzler, C., Schniter, P., Veeraraghavan, A., and Baraniuk, R. prDeep: Robust phase retrieval with a flexible deep network. In *Proc. 36th Int. Conf. Mach. Learn.*, pp. 3501–3510, Jul. 10–15 2018.
- Milanfar, P. and Delbracio, M. Denoising: A powerful building-block for imaging, inverse problems, and machine learning. *arXiv preprint arXiv:2409.06219*, 2024.
- Millard, C. and Chiew, M. A theoretical framework for self-supervised mr image reconstruction using sub-sampling via variable density noisier2noise. *IEEE Trans. Comp. Imag.*, 2023.
- Monga, V., Li, Y., and Eldar, Y. C. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Process. Mag.*, 38(2):18–44, March 2021.

- Nathan, O., Levy, D., Treibitz, T., and Rosenbaum, D. Osmosis: Rgb-d diffusion prior for underwater image restoration. In *Euro. Conf. Comp. Vis. (ICCV)*, pp. 302–319. Springer, 2024.
- Nesterov, Y. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer Academic Publishers, 2004.
- Ongie, G., Jalal, A., Metzler, C. A., Baraniuk, R. G., Dimakis, A. G., and Willett, R. Deep learning techniques for inverse problems in imaging. *IEEE J. Sel. Areas Inf. Theory*, 1(1):39–56, May 2020.
- Quan, Y., Chen, M., Pang, T., and Ji, H. Self2self with dropout: Learning self-supervised denoising from single image. In *Proc. IEEE Conf. Comput. Vis. and Pattern Recognit. (CVPR)*, pp. 1890–1898, 2020.
- Reehorst, E. T. and Schniter, P. Regularization by denoising: Clarifications and new interpretations. *IEEE Trans. Comput. Imag.*, 5(1):52–67, March 2019.
- Renaud, M., Prost, J., Leclaire, A., and Papadakis, N. Plug-and-play image restoration with stochastic denoising regularization. In *Int. Conf. Machine Learning (ICML)*, 2024.
- Romano, Y., Elad, M., and Milanfar, P. The little engine that could: Regularization by denoising (RED). *SIAM J. Imaging Sci.*, 10(4):1804–1844, 2017.
- Rout, L., Chen, Y., Kumar, A., Caramanis, C., Shakkottai, S., and Chu, W. Beyond first-order tweedie: Solving inverse problems using latent diffusion. In *Proc. IEEE Conf. Comput. Vis. and Pattern Recognit. (CVPR)*, pp. 9472–9481, 2024.
- Ryu, E. K., Liu, J., Wang, S., Chen, X., Wang, Z., and Yin, W. Plug-and-play methods provably converge with properly trained denoisers. In *Proc. Int. Conf. Machine Learning (ICML)*, volume 97, pp. 5546–5557, 2019.
- Song, B., Kwon, S. M., Zhang, Z., Hu, X., Qu, Q., and Shen, L. Solving inverse problems with latent diffusion models via hard data consistency. In *Proc. Int. Conf. on Learn. Represent. (ICLR)*, 2024.
- Song, Y., S., J., K., D. P., K., A., E., S., and P., B. Score-based generative modeling through stochastic differential equations. In *Proc. Int. Conf. on Learn. Represent. (ICLR)*, 2021.
- Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C. L., Yakubova, N., Knoll, F., and Johnson, P. M. End-to-end variational networks for accelerated MRI reconstruction. *arXiv*, 2004.06688, 2020.
- Sun, Y., Wohlberg, B., and Kamilov, U. S. An online plug-and-play algorithm for regularized image reconstruction. *IEEE Trans. Comput. Imag.*, 5(3):395–408, September 2019.
- Sun, Y., Xu, S., Li, Y., Tian, L., Wohlberg, B., and Kamilov, U. S. Regularized Fourier ptychography using an online plug-and-play algorithm. In *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Process. (ICASSP)*, pp. 7665–7669, Brighton, UK, May 12–17, 2019.
- Sun, Y., Wu, Z., Wohlberg, B., and Kamilov, U. S. Scalable plug-and-play ADMM with convergence guarantees. *IEEE Trans. Comput. Imag.*, 7:849–863, July 2021.
- Sun, Y., Wu, Z., Chen, Y., Feng, B. T., and Bouman, K. L. Provable probabilistic imaging using score-based generative priors. *IEEE Trans. Comp. Imag.*, 2024.
- Tachella, J., Dongdong, C., and Davies, M. Unsupervised learning from incomplete measurements for inverse problems. In *Proc. Adv. Neural Inf. Process. Syst.*, pp. 4983–4995, 2022.
- Teodoro, A. M., Bioucas-Dias, J. M., and Figueiredo, M. A convergent image fusion algorithm using scene-adapted Gaussian-mixture-based denoising. *IEEE Trans. Image Process.*, 28(1):451–463, January 2019.
- Terris, M. and Moreau, T. Meta-prior: Meta learning for adaptive inverse problem solvers. *arXiv preprint arXiv:2311.18710*, 2023.
- Terris, M., Moreau, T., Pustelnik, N., and Tachella, J. Equivariant plug-and-play image reconstruction. In *Proc. IEEE Conf. Comput. Vis. and Pattern Recognit. (CVPR)*, pp. 25255–25264, 2024.
- Tirer, T. and Giryes, R. Image restoration by iterative denoising and backward projections. *IEEE Trans. Image Process.*, 28(3):1220–1234, Mar. 2019.
- Uecker, M., Lai, P., Murphy, M. J., Virtue, P., Elad, M., Pauly, J. M., Vasanawala, S. S., and Lustig, M. ESPIRiT— an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magn. Reson. Med.*, 71(3):990–1001, 2014.
- Ulyanov, D., Vedaldi, A., and Lempitsky, V. Deep image prior. In *Proc. IEEE Conf. Comput. Vis. and Pattern Recognit. (CVPR)*, pp. 9446–9454, 2018.
- Venkatakrishnan, S. V., Bouman, C. A., and Wohlberg, B. Plug-and-play priors for model based reconstruction. In *Proc. IEEE Global Conf. Signal Process. and Inf. Process.*, pp. 945–948, Dec. 3–5, 2013.
- Wang, S., Su, Z., Ying, L., Peng, X., Zhu, S., Liang, F., Feng, D., and Liang, D. Accelerating magnetic resonance imaging via deep learning. In *Proc. Int. Symp. Biomedical*

- Imaging (ISBI)*, pp. 514–517, April 2016. doi: 10.1109/ISBI.2016.7493320.
- Wang, Y., Yu, J., and Zhang, J. Zero-shot image restoration using denoising diffusion null-space model. In *Proc. Int. Conf. on Learn. Represent. (ICLR)*, 2023.
- Wei, K., Aviles-Rivero, A., Liang, J., Fu, Y., Schönlieb, C.-B., and Huang, H. Tuning-free plug-and-play proximal algorithm for inverse imaging problems. In *Int. Conf. Machine Learning (ICML)*, 2020.
- Wen, B., Ravishankar, S., Zhao, Z., Giryes, R., and Ye, J. C. Physics-driven machine learning for computational imaging. *IEEE Signal Process. Mag.*, 40(1):28–30, 2023.
- Wu, Z., Sun, Y., Chen, Y., Zhang, B., Yue, Y., and Bouman, K. L. Principled probabilistic imaging using diffusion models as plug-and-play priors. *arXiv preprint arXiv:2405.18782*, 2024.
- Xia, B., Zhang, Y., Wang, S., Wang, Y., Wu, X., Tian, Y., Yang, W., and Van Gool, L. Diffir: Efficient diffusion model for image restoration. In *Proc. IEEE Int. Conf. Comp. Vis. (ICCV)*, pp. 13095–13105, 2023.
- Xiao, J., Feng, R., Zhang, H., Liu, Z., Yang, Z., Zhu, Y., Fu, X., Zhu, K., Liu, Y., and Zha, Z. Dreamclean: Restoring clean image using deep diffusion prior. In *Proc. Int. Conf. on Learn. Represent. (ICLR)*, 2024.
- Yaman, B., Hosseini, S. A. H., Moeller, S., Ellermann, J., Uğurbil, K., and Akçakaya, M. Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data. *Magn. Reson. Med.*, 84(6): 3172–3191, Jul. 2020.
- Zhang, J. and Ghanem, B. ISTA-Net: Interpretable optimization-inspired deep network for image compressive sensing. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 1828–1837, 2018.
- Zhang, K., Zuo, W., and Zhang, L. Deep plug-and-play super-resolution for arbitrary blur kernels. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 1671–1681, Long Beach, CA, USA, June 16-20, 2019.
- Zhang, K., Li, Y., Zuo, W., Zhang, L., Van Gool, L., and Timofte, R. Plug-and-play image restoration with deep denoiser prior. *IEEE Trans. Patt. Anal. and Machine Intell.*, 44(10):6360–6376, October 2022.
- Zhao, J., Song, B., and Shen, L. Cosign: Few-step guidance of consistency model to solve general inverse problems. *arXiv preprint arXiv:2407.12676*, 2024.
- Zhu, Y., Zhang, K., Liang, J., Cao, J., Wen, B., Timofte, R., and Van G., L. Denoising diffusion models for plug-and-play image restoration. In *Proc. IEEE Conf. Comput. Vis. and Pattern Recognit. (CVPR)*, pp. 1219–1229, 2023.

Here’s what we cover in the appendix:

- **Theoretical Analysis of ShaRP**

- Proof of Theorem 1 (detailed in Subsection A.1): Derivation of the gradient of the ShaRP regularizer.
- Proof of Proposition 1 (detailed in Subsection A.2): Convergence guarantees for ShaRP with inexact Minimum Mean Square Error (MMSE) restoration operators using biased Stochastic Gradient Descent (SGD) analysis.
- Proof of Theorem 2 (detailed in Subsection A.3): Equivalence of the MMSE estimator learned via weighted self-supervised loss to its supervised counterpart.

- **Experiment Details** (Section B)

- Implementation details for Compressed Sensing MRI (CS-MRI) tasks (Subsection B.1):
- Implementation details for Single Image Super-Resolution (SISR) tasks (Subsection B.2):

- **Additional results for CS-MRI** (Section C)

- Quantitative performance of ShaRP for uniform (Table 5) and random (Table 6) subsampling settings.
- Visual comparisons for CS-MRI tasks (Figure 5, Figure 6).
- Performance of ShaRP with self-supervised restoration priors (Subsection C.2, Table 7).
- Convergence performance analysis of ShaRP with both supervised and self-supervised priors (Subsection C.3, Figure 7), aligning with the convergence analysis for inexact MMSE estimators.
- Performance comparison with additional baseline methods on matched and mismatched settings (Table 8, Figure 8).

- **Additional visual results for SISR** (Section D)

- Further visual comparisons against various baseline methods for SISR (Figure 9, Figure 10).
- Additional visual comparisons against the DRP method (Figure 11).
- Quantitative (Table 10) and visual (Figure 12) comparisons against additional diffusion-based baselines.

- **Additional Experiments**

- Ablation study on using a pre-trained super-resolution prior for the CS-MRI task (Subsection E.1, Algorithm 5, Table 12).
- Ablation study on the impact of the number of restoration priors,  $b$ , in the ensemble on CS-MRI performance (Subsection E.2, Figure 13, Figure 14).
- Ablation study on the influence of the hyperparameter  $\alpha$  (controlling prior selection, as introduced in Subsection B.1) on CS-MRI performance (Subsection E.3, Figure 15).

## A. Theoretical Analysis of ShaRP

### A.1. Proof of Theorem 1

**Theorem.** Assume that the prior density  $p_{\mathbf{x}}$  is non-degenerate over  $\mathbb{R}^n$  and let  $R^*$  be the MMSE restoration operator (4) corresponding to the restoration problems (3). Then, we have that

$$\nabla h(\mathbf{x}) = \frac{\tau}{\sigma^2} \left( \mathbb{E}_{\mathbf{s} \sim G_{\sigma}(\mathbf{s} - \mathbf{H}\mathbf{x}), \mathbf{H} \sim p_{\mathbf{H}}} \left[ \mathbf{H}^{\top} \mathbf{H} (\mathbf{x} - R^*(\mathbf{s}, \mathbf{H})) \right] \right),$$

where  $h$  is the ShaRP regularizer in (6).

*Proof.* The ShaRP regularizer  $h(\mathbf{x})$  is defined as

$$\begin{aligned} h(\mathbf{x}) &= \tau \mathbb{E}_{\mathbf{s} \sim G_{\sigma}(\mathbf{s} - \mathbf{H}\mathbf{x}), \mathbf{H} \sim p_{\mathbf{H}}} \left[ -\log p(\mathbf{s} | \mathbf{H}) \right] \\ &= -\tau \int p(\mathbf{H}) \left[ \int G_{\sigma}(\mathbf{s} - \mathbf{H}\mathbf{x}) \log p(\mathbf{s} | \mathbf{H}) \, d\mathbf{s} \right] d\mathbf{H}, \end{aligned} \quad (13)$$

where  $G_\sigma$  is the Gaussian probability density with variance  $\sigma^2$  and  $p(\mathbf{s}|\mathbf{H})$  is the likelihood function for the degraded observation given the operator  $\mathbf{H}$ . The expectation over  $p(\mathbf{H})$  accounts for the randomness of the restoration operator  $\mathbf{H}$ .

We start by relating the MMSE restoration operator to the score of the degraded observation

$$\nabla p(\mathbf{s}|\mathbf{H}) = \frac{1}{\sigma^2} \int (\mathbf{H}\mathbf{x} - \mathbf{s}) G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x}) p_{\mathbf{x}}(\mathbf{x}) d\mathbf{x},$$

where  $p_{\mathbf{x}}$  is the prior. By using the definition of the MMSE estimator, we obtain the relationship

$$\nabla \log p(\mathbf{s}|\mathbf{H}) = \frac{1}{\sigma^2} (\mathbf{H}\mathbf{R}^*(\mathbf{s}, \mathbf{H}) - \mathbf{s}). \quad (14)$$

Consider the function inside the parenthesis in the expression for the ShaRP regularizer (13)

$$\rho(\mathbf{z}) := (G_\sigma * \log p_{\mathbf{s}|\mathbf{H}})(\mathbf{z}) = \int G_\sigma(\mathbf{z} - \mathbf{s}) \log p(\mathbf{s}|\mathbf{H}) d\mathbf{s},$$

where  $\mathbf{z}$  has the same dimensions as  $\mathbf{s}$  and  $*$  denotes convolution. The gradient of  $\rho$  is given by

$$\begin{aligned} \nabla \rho(\mathbf{z}) &= (\nabla G_\sigma * \log p_{\mathbf{s}|\mathbf{H}})(\mathbf{z}) = (G_\sigma * \nabla \log p_{\mathbf{s}|\mathbf{H}})(\mathbf{z}) \\ &= \frac{1}{\sigma^2} \int G_\sigma(\mathbf{z} - \mathbf{s}) [\mathbf{H}\mathbf{R}^*(\mathbf{s}, \mathbf{H}) - \mathbf{s}] d\mathbf{s} \\ &= \frac{1}{\sigma^2} \left( \mathbf{H} \int \mathbf{R}^*(\mathbf{s}, \mathbf{H}) G_\sigma(\mathbf{z} - \mathbf{s}) d\mathbf{s} - \mathbf{z} \right) \end{aligned}$$

where we used (14). By using  $\mathbf{z} = \mathbf{H}\mathbf{x}$ , we write the gradient with respect to  $\mathbf{x}$

$$\nabla_{\mathbf{x}} \rho(\mathbf{H}\mathbf{x}) = \frac{1}{\sigma^2} \mathbf{H}^\top \mathbf{H} \left( \int \mathbf{R}^*(\mathbf{s}, \mathbf{H}) G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x}) d\mathbf{s} - \mathbf{x} \right)$$

By using this expression in (13), we obtain the desired result

$$\begin{aligned} \nabla h(\mathbf{x}) &= -\frac{\tau}{\sigma^2} \left[ \int p(\mathbf{H}) \int G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x}) (\mathbf{H}^\top \mathbf{H}(\mathbf{R}^*(\mathbf{s}, \mathbf{H}) - \mathbf{x})) d\mathbf{s} d\mathbf{H} \right] \\ &= \frac{\tau}{\sigma^2} \mathbb{E}_{\mathbf{s} \sim G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x}), \mathbf{H} \sim p_{\mathbf{H}}} [\mathbf{H}^\top \mathbf{H}(\mathbf{x} - \mathbf{R}^*(\mathbf{s}, \mathbf{H}))]. \end{aligned}$$

□

## A.2. Proof of Proposition 1

We adopt the analysis of biased stochastic gradient descent (SGD) from the optimization literature (Bertsekas, 2011; Ghadimi & Lan, 2016; Demidovich et al., 2023) to provide theoretical convergence guarantees to ShaRP under inexact MMSE restoration operators. Our contribution here is the interpretation of the iterations of ShaRP as a variant of biased SGD, rather than a general analysis of SGD, which is well-known in the optimization literature.

**Proposition.** *Run ShaRP for  $t \geq 1$  iterations using the step-size  $0 < \gamma \leq 1/L$  under Assumptions 1-3. Then, the sequence  $\mathbf{x}^k$  generated by ShaRP satisfies*

$$\mathbb{E} \left[ \frac{1}{t} \sum_{k=1}^t \|\nabla f(\mathbf{x}^{k-1})\|_2^2 \right] \leq \frac{2}{t} (f(\mathbf{x}^0) - f^*) + \gamma L \nu^2 + \varepsilon^2.$$

*Proof.* First note that from the definition of the bias in eq. (9), we have that

$$\mathbb{E} \left[ \widehat{\nabla} f(\mathbf{x}^{k-1}) | \mathbf{x}^{k-1} \right] = \nabla f(\mathbf{x}^{k-1}) + \mathbf{b}(\mathbf{x}^{k-1}), \quad (15)$$

where the expectation is with respect to  $\mathbf{s} \sim G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x}^{k-1})$  and  $\mathbf{H} \sim p_{\mathbf{H}}$ . In order to simplify the notation, we will drop these subscripts from the expectations in the analysis below.

Consider the iteration  $k \geq 1$  of ShaRP with inexact MMSE operator

$$\begin{aligned} f(\mathbf{x}^k) &\leq f(\mathbf{x}^{k-1}) + \nabla f(\mathbf{x}^{k-1})^\top (\mathbf{x}^k - \mathbf{x}^{k-1}) + \frac{L}{2} \|\mathbf{x}^k - \mathbf{x}^{k-1}\|_2^2 \\ &= f(\mathbf{x}^{k-1}) - \gamma \nabla f(\mathbf{x}^{k-1})^\top \widehat{\nabla} f(\mathbf{x}^{k-1}) + \frac{L\gamma^2}{2} \|\widehat{\nabla} f(\mathbf{x}^{k-1})\|_2^2, \end{aligned}$$

where in the first line we used the Lipschitz continuity of  $\nabla f$ . By taking the expectation with respect to  $\mathbf{s} \sim G_\sigma(\mathbf{s} - \mathbf{H}\mathbf{x}^{k-1})$  and  $\mathbf{H} \sim p_{\mathbf{H}}$  on both sides of this expression, we get

$$\begin{aligned} \mathbb{E}[f(\mathbf{x}^k)|\mathbf{x}^{k-1}] &\leq f(\mathbf{x}^{k-1}) - \gamma \nabla f(\mathbf{x}^{k-1})^\top (\nabla f(\mathbf{x}^{k-1}) + \mathbf{b}(\mathbf{x}^{k-1})) + \frac{L\gamma^2}{2} \mathbb{E} \left[ \|\widehat{\nabla} f(\mathbf{x}^{k-1})\|_2^2 | \mathbf{x}^{k-1} \right] \\ &\leq f(\mathbf{x}^{k-1}) - \frac{\gamma}{2} \|\nabla f(\mathbf{x}^{k-1})\|_2^2 + \frac{\gamma}{2} \|\mathbf{b}(\mathbf{x}^{k-1})\|_2^2 \\ &\quad + \frac{L\gamma^2}{2} \left( \mathbb{E} \left[ \|\widehat{\nabla} f(\mathbf{x}^{k-1})\|_2^2 | \mathbf{x}^{k-1} \right] - \left( \mathbb{E}[\widehat{\nabla} f(\mathbf{x}^{k-1}) | \mathbf{x}^{k-1}] \right)^2 \right) \\ &\leq f(\mathbf{x}^{k-1}) - \frac{\gamma}{2} \|\nabla f(\mathbf{x}^{k-1})\|_2^2 + \frac{\gamma\varepsilon^2}{2} + \frac{L\gamma^2\nu^2}{2}. \end{aligned}$$

In the second row, we completed the square, applied eq. (15), and used the assumption that  $\gamma \leq 1/L$ . In the third row, we used the variance and bias bounds in Assumptions 2 and 3. By rearranging the expression, we get the following bound

$$\|\nabla f(\mathbf{x}^{k-1})\|_2^2 \leq \frac{2}{\gamma} (f(\mathbf{x}^{k-1}) - \mathbb{E}[f(\mathbf{x}^k)|\mathbf{x}^{k-1}]) + L\gamma\nu^2 + \varepsilon^2$$

By taking the total expectation, averaging over  $t$  iterations, and using the lower bound  $f^*$ , we get the desired result

$$\mathbb{E} \left[ \frac{1}{t} \sum_{k=1}^t \|\nabla f(\mathbf{x}^{k-1})\|_2^2 \right] \leq \frac{2}{\gamma t} (f(\mathbf{x}^0) - f^*) + L\gamma\nu^2 + \varepsilon^2.$$

□

### A.3. Proof of Theorem 2:

Let the undersampled measurements be defined as  $\mathbf{s} = \mathbf{H}\mathbf{x} + \mathbf{n}$ , where  $\mathbf{H} = \mathbf{P}\mathbf{M}$ . Here,  $\mathbf{P}$  represents a binary subsampling matrix,  $\mathbf{M}$  denotes the square transfer operator, and  $\mathbf{n}$  corresponds to the noise vector.

To show that an MMSE estimator can be learned from undersampled measurement, we need the following assumption.

**Assumption 5.**  $\mathbb{E}_{\mathbf{P}}[\mathbf{P}^\top \mathbf{P}]$  has a full rank and  $\mathbf{M}$  is an orthogonal matrix, where the expectation is taken over  $p_{\mathbf{P}}$ .

This assumption implies that the *union* of all sampling matrices  $\mathbf{P}$  spans the complete measurement domain, even though each individual  $\mathbf{P}$  may remain undersampled. This property can be achieved by incorporating an additional weight  $\mathbf{W}$  into the loss function, where:  $\mathbf{W} = \mathbf{P}'\overline{\mathbf{W}}(\mathbf{P}'\overline{\mathbf{W}})^\top \in \mathbb{R}^{m \times m}$  denotes a subsampled variant of  $\overline{\mathbf{W}}$  given  $\mathbf{P}'$ .

**Proposition 2.** *When Assumption 2 is satisfied,*

$$\mathbb{E}[\mathbf{M}^\top \mathbf{P}'^\top \mathbf{W} \mathbf{P}' \mathbf{M}] = \mathbf{I},$$

where the expectation is with respect to  $p_{\mathbf{M}}$ . This proof is the same as provided in previous work (Gan et al., 2023b).

**Theorem 3.** *Under Assumption 5, the MMSE estimator  $\mathbf{R}$  learned using the weighted self-supervised loss ( $\ell_{\text{self}}$ ) is equivalent to its supervised counterpart ( $\ell_{\text{sup}}$ ). Specifically, we have:*

$$\mathbf{R}_{\ell_{\text{self}}}(\boldsymbol{\theta}) = \mathbf{R}_{\ell_{\text{sup}}}(\boldsymbol{\theta}). \quad (16)$$

where

$$\ell_{\text{sup}} = \mathbb{E} \left[ \frac{1}{2} \|\bar{\mathbf{x}} - \mathbf{x}\|_2^2 \right] \quad (17)$$

and

$$\ell_{\text{self}} = \mathbb{E} \left[ \frac{1}{2} \|\mathbf{P}'\mathbf{M}\bar{\mathbf{x}} - \mathbf{s}'\|_{\mathbf{W}}^2 \right]. \quad (18)$$

The vector  $\bar{\mathbf{x}} = \mathbf{R}(\mathbf{s})$  is MMSE estimation of  $\mathbf{R}$  for  $\mathbf{s}$ .

*Proof.* Note that the self-supervised loss involves the term  $\mathbf{P}'\mathbf{M}\bar{\mathbf{x}} - \mathbf{s}'$ , where  $\mathbf{s}' = \mathbf{P}'\mathbf{M}\mathbf{x} + \mathbf{n}'$

$$\mathbf{P}'\mathbf{M}\bar{\mathbf{x}} - \mathbf{s}' = \mathbf{P}'\mathbf{M}(\bar{\mathbf{x}} - \mathbf{x}) - \mathbf{n}'. \quad (19)$$

Thus, the self-supervised loss becomes:

$$\ell_{\text{self}} = \mathbb{E} \left[ \frac{1}{2} \|\mathbf{P}'\mathbf{M}(\bar{\mathbf{x}} - \mathbf{x}) - \mathbf{n}'\|_{\mathbf{W}}^2 \right]. \quad (20)$$

Expanding the squared term:

$$\begin{aligned} \|\mathbf{P}'\mathbf{M}(\bar{\mathbf{x}} - \mathbf{x}) - \mathbf{n}'\|_{\mathbf{W}}^2 &= \|\mathbf{P}'\mathbf{M}(\bar{\mathbf{x}} - \mathbf{x})\|_{\mathbf{W}}^2 - 2(\mathbf{P}'\mathbf{M}(\bar{\mathbf{x}} - \mathbf{x}))^\top \mathbf{W}\mathbf{n}' + \|\mathbf{n}'\|_{\mathbf{W}}^2 \\ &= (\bar{\mathbf{x}} - \mathbf{x})^\top \mathbf{M}^\top \mathbf{P}'^\top \mathbf{W}\mathbf{P}'\mathbf{M}(\bar{\mathbf{x}} - \mathbf{x}) - 2(\mathbf{P}'\mathbf{M}(\bar{\mathbf{x}} - \mathbf{x}))^\top \mathbf{W}\mathbf{n}' + \|\mathbf{n}'\|_{\mathbf{W}}^2. \end{aligned}$$

So that

$$\begin{aligned} &\mathbb{E} [\|\mathbf{P}'\mathbf{M}(\bar{\mathbf{x}} - \mathbf{x}) - \mathbf{n}'\|_{\mathbf{W}}^2] \\ &= \mathbb{E} \left[ (\bar{\mathbf{x}} - \mathbf{x})^\top \mathbf{M}^\top \mathbf{P}'^\top \mathbf{W}\mathbf{P}'\mathbf{M}(\bar{\mathbf{x}} - \mathbf{x}) \right] - \mathbb{E} [2(\mathbf{P}'\mathbf{M}(\bar{\mathbf{x}} - \mathbf{x}))^\top \mathbf{W}\mathbf{n}'] + \mathbb{E} [\|\mathbf{n}'\|_{\mathbf{W}}^2]. \\ &= \mathbb{E} [\|\bar{\mathbf{x}} - \mathbf{x}\|_2^2] + \text{constant}, \end{aligned}$$

where the first term equals to  $\mathbb{E} [\|\bar{\mathbf{x}} - \mathbf{x}\|_2^2]$  due to the Proposition 2 that  $\mathbb{E} [\mathbf{M}^\top \mathbf{P}'^\top \mathbf{W}\mathbf{P}'\mathbf{M}] = \mathbf{I}$ ; the second term equals to zero because  $\mathbf{n}'$  is zero-mean and independent of  $\mathbf{P}'$  and  $\mathbf{x}$ ; The third term,  $\|\mathbf{n}'\|_{\mathbf{W}}^2$ , does not depend on  $\mathbf{x}$  and contributes a constant that does not affect the optimization for training the MMSE estimator  $\mathbf{R}$ .  $\square$

## B. Experiment Details

### B.1. Implementation details of CS-MRI tasks

**Dataset.** We simulated multi-coil subsampled measurements using T2-weighted human brain MRI data from the open-access fastMRI dataset, which comprises 4,912 fully sampled multi-coil slices for training and 470 slices for testing. Each slice has been cropped into a complex-valued image with dimensions  $320 \times 320$ . The coil sensitivity maps for each slice are precomputed using the ESPIRiT algorithm (Uecker et al., 2014). We simulated a Cartesian sampling pattern that subsamples along the  $k_y$  dimension while fully sampling along the  $k_x$  dimension.

**Subsampling pattern for CS-MRI.** In this paper, we explored two types of subsampling patterns for MRI reconstruction tasks. All undersampling masks were generated by first including a set number of *auto-calibration signal (ACS)* lines, ensuring a fully-sampled central k-space region.

Figure 4 illustrates the k-space trajectories for both random and uniform (equidistant) subsampling at acceleration factors of 4, 6, and 8. Notably, different patterns were used for training and testing. During training, our restoration prior was only trained on a uniform mask with a subsampling rate of 6. However, for inference, we employed both uniform and random masks at subsampling rates of 4 and 6, creating a mismatch between the pre-trained restoration prior and the test configurations.

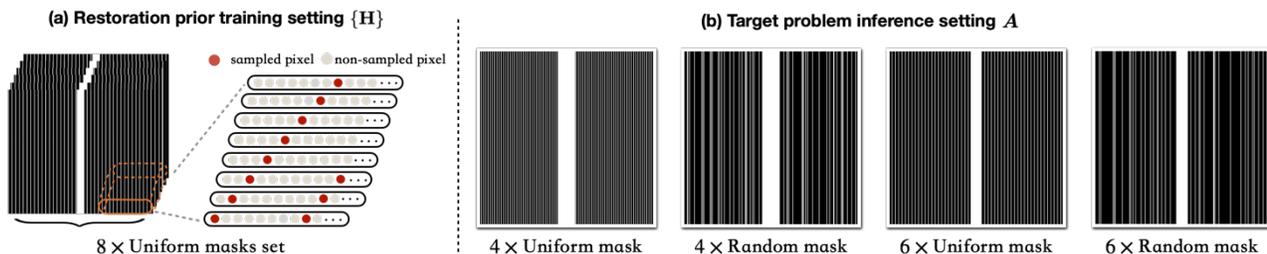


Figure 4: Illustration of the undersampling masks used for CS-MRI in this work. (a) The eight different  $8 \times$  uniform masks used for training the restoration prior. (b) The inference setting for ShaRP, demonstrating how the prior trained on the masks in (a) can be applied to solve other problems without retraining.

---

#### Algorithm 2 Supervised Training of CS-MRI Restoration Network

---

**Require:** dataset:  $p(x)$ , sampling operator set:  $\{M_1, M_2, \dots, M_8\}$ , Restoration model:  $R_\theta(\cdot, \alpha)$

**repeat:**

$$x \sim p(x), M \sim \{M_1, M_2, \dots, M_8\}, e \sim \mathcal{N}(0, \sigma^2 \mathbf{I}), \alpha \sim \mathcal{U}([0, 1])$$

$$y = Mx + e$$

$$\min_\theta \|R_\theta((1 - \alpha)x + \alpha M^T y; \alpha) - x\|_2^2$$

**until converge**

---

#### B.1.1. IMPLEMENTATION OF SUPERVISED PRIOR FOR CS-MRI

**Models training for supervised case.** We use the same U-Net architecture as employed in the official implementation of DDS<sup>2</sup> for  $R(\cdot; \alpha)$ . For the supervised learning case, we select 1,000 different  $\alpha$  values to train the model, following the  $\alpha$  schedule outlined by I<sup>2</sup>SB (Liu et al., 2023). The model is trained with Adam optimizer with a learning rate of  $5 \times 10^{-5}$ . As shown in Algorithm 2, we train our supervised learning model using eight different masks for  $8 \times$  uniform sampling CS-MRI reconstruction. In the pseudocode,  $\{M_1, M_2, \dots, M_8\}$  represent the eight different MRI degradation operators, each defined by a unique sampling pattern, as shown in Figure 4 (a). This results in a total of 8,000 possible combinations of  $\alpha$  values and sampling masks, effectively creating an ensemble of restoration priors during training.

**Inference with a Subset of the Ensemble (Supervised Case).** During inference, to simplify computation and focus on the most effective priors, we use only a subset of the supervised trained ensemble. Specifically, we fix the  $\alpha$  value to a particular choice (e.g.,  $\alpha = 0.5$ ) and use the 8 different sampling masks  $\{M_1, M_2, \dots, M_8\}$ , resulting in 8 restoration priors.

<sup>2</sup><https://github.com/HJ-harry/DDS>

**Step size and regularization parameter.** To ensure fairness, for each problem setting, each method—both proposed and baseline—is fine-tuned for optimal PSNR using 10 slices from a validation set separate from the test set. The same step size  $\gamma$  and regularization parameter  $\tau$  are then applied consistently across the entire test set.

**Baseline details.** We compare ShaRP with several variants of denoiser- and diffusion model-based methods. For denoiser-based approaches, we include PnP-FISTA (Kamilov et al., 2023), PnP-ADMM (Chan et al., 2017). PnP-FISTA and PnP-ADMM correspond to the FISTA and ADMM variants of PnP, both utilizing AWGN denoisers built on DRUNet (Zhang et al., 2022). For diffusion model-based methods, we compare with DPS (Chung et al., 2023) and DDS (Chung et al., 2024), which use pre-trained diffusion models as priors and apply different posterior sampling strategies to address general inverse problems. We use the same pre-trained diffusion model configuration as outlined in the DDS paper. For all baseline methods, we fine-tuned their parameters to maximize the PSNR value. Notably, both the DRUNet denoiser and the diffusion model were trained using the same dataset employed for training our restoration prior. For a fair comparison, the diffusion model pre-trained for DDS and DPS use the same network architecture as our restoration network. All models are trained from scratch on the fastMRI training set, following the architecture settings provided in DDS<sup>3</sup>. We also compared with method that also use the deep restoration prior to solve general inverse problem: DRP (Hu et al., 2024c). For DRP, we utilize the same pre-trained restoration network as in ShaRP. However, instead of employing a set of degradation priors, DRP uses a single fixed prior. For a fair comparison, we selected the optimal fixed prior—defined by a fixed  $\alpha$  and subsampling mask—based on PSNR performance on the validation set, and applied it accordingly.

B.1.2. IMPLEMENTATION OF SELF-SUPERVISED PRIOR FOR CS-MRI

---

**Algorithm 3** Self-Supervised Training of CS-MRI Restoration Network

---

**Require:** dataset:  $p(\mathbf{y}_i, \mathbf{M}_i, \mathbf{y}_j, \mathbf{M}_j)$ , Restoration model:  $R_\theta(\cdot)$

**repeat:**

$\mathbf{y}_i, \mathbf{M}_i, \mathbf{y}_j, \mathbf{M}_j \sim p(\mathbf{y}_i, \mathbf{M}_i, \mathbf{y}_j, \mathbf{M}_j), e \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$

$\min_\theta \| \mathbf{M}_j R_\theta(\mathbf{M}_i^\top (\mathbf{y}_i + e)) - \mathbf{y}_j \|_{\mathbf{W}}^2$

**until converge**

---

**Models training for (Self-Supervised Case).** For self-supervised training, the ground truth reference  $\mathbf{x}$  is not available as a label. Instead, as shown in Algorithm 3, we work with pairs of subsampled measurements,  $\mathbf{y}_i$  and  $\mathbf{y}_j$ , along with their corresponding sampling operators,  $\mathbf{M}_i$  and  $\mathbf{M}_j$ . These paired measurements exhibit significant overlap within the shared *auto-calibration signal (ACS)* region, which increases the weighting of these overlapping k-space regions. Following the approach proposed by SSDEQ (Gan et al., 2023b), we introduce a diagonal weighting matrix  $\mathbf{W}$  to account for the oversampled regions in the loss function. By incorporating this weighted loss, we are able to train our MMSE restoration operator using incomplete measurements alone. Furthermore, unlike the supervised case where we use the combination of  $\alpha$  values to form an ensemble, in the self-supervised setting, we construct the ensemble using only eight different sampling masks across the entire dataset.

**Inference Using All Restoration Priors (Self-Supervised Case).** During inference in the self-supervised setting, we utilize all 8 restoration priors corresponding to the different sampling masks. By incorporating the entire ensemble, we fully leverage its capacity to remove the artifacts and enhance reconstruction performance.

**Step size and regularization parameter.** To ensure fairness, for each problem setting, each method—both proposed and baseline—is fine-tuned for optimal PSNR using 10 slices from a validation set separate from the test set. The same step size  $\gamma$  and regularization parameter  $\tau$  are then applied consistently across the entire test set.

**Baseline details.** In the self-supervised setting, we compared ShaRP with two widely used traditional methods: TV (Block et al., 2007) and GRAPPA (Griswold et al., 2002), both of which address the restoration problem without requiring fully-sampled references. Additionally, we included SPICER (Hu et al., 2024d), a recent state-of-the-art self-supervised deep unrolling method designed for MRI reconstruction using only pairs of undersampled measurements. To ensure consistency, we trained the SPICER model on the same amount of paired data used for training our restoration prior in the  $8 \times$  uniform CS-MRI setting and applied it to other CS-MRI configurations.

<sup>3</sup><https://github.com/HJ-harry/DDS>

**Algorithm 4** Gaussian Deblurring Restoration network training

**Require:** dataset:  $p(\mathbf{x}, \mathbf{y})$ , Gaussian blur operator:  $\mathbf{K}$ ,  $R_\theta(\cdot, \alpha)$   
**repeat:**  
 $\mathbf{x} \sim p(\mathbf{x})$ ,  $\mathbf{e} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ ,  $\alpha \sim \mathcal{U}([0, 1])$   
 $\min_\theta \|R_\theta((1 - \alpha)\mathbf{x} + \alpha\mathbf{K}\mathbf{x}; \alpha) - \mathbf{x}\|_2^2$   
**until converge**

**B.2. Implementation details of SISR tasks**

**Restoration Model training.** We use the same U-Net architecture as the Gaussian deblurring model provided by I<sup>2</sup>SB<sup>4</sup>. Utilizing the pre-trained checkpoints from their repository, we fine-tune our model accordingly. Specifically, we align with their codebase and configure the model type to OT-ODE to satisfy our MMSE restoration operator assumption.

To create an ensemble of restoration priors, we consider a family of degradation operators that are convex combinations of the identity mapping  $\mathbf{I}$  and the Gaussian blur operator  $\mathbf{K}$ . The blurring operator  $\mathbf{K}$  corresponds to convolution with a Gaussian blur kernel of size  $31 \times 31$  and standard deviation 3. Specifically, we define the degradation operator as  $\mathbf{H}_\alpha = (1 - \alpha)\mathbf{I} + \alpha\mathbf{K}$ , where  $\alpha \in [0, 1]$  controls the degradation level. By varying  $\alpha$ , we generate multiple degradation operators, allowing us to train the restoration network  $R$  to handle all these operators, expressed as  $R(\mathbf{s}, \mathbf{H}_\alpha) = \mathbb{E}[\mathbf{x}|\mathbf{s}, \mathbf{H}_\alpha]$ , where  $\mathbf{s}$  is the degraded image and  $\mathbf{x}$  is the original image.

We select 1,000 different  $\alpha$  values from the interval  $[0, 1]$ , following the  $\alpha$  schedule outlined by I<sup>2</sup>SB (Liu et al., 2023). This results in 1,000 different degradation operators  $\mathbf{H}_\alpha$ , effectively creating an ensemble of restoration priors during training. The model is trained using the Adam optimizer with a learning rate of  $5 \times 10^{-5}$ .

**Inference with a Subset of the Ensemble.** During inference, to simplify computation and focus on the most effective priors, we use only a subset of the supervised trained ensemble. Specifically, we select 6  $\alpha$  values, resulting in 6 restoration priors.

**Step size and regularization parameter.** To ensure fairness, for each problem setting, each method—both proposed and baseline—is fine-tuned for optimal PSNR using 5 images from a validation set separate from the test set. The same step size  $\gamma$  and regularization parameter  $\tau$  are then applied consistently across the entire test set.

**Baseline details.** We compare ShaRP against several denoiser- and diffusion model-based methods. For denoiser-based approaches, we evaluate DPIP (Zhang et al., 2022), which relies on half-quadratic splitting (HQS) iterations with DRUNet denoisers. For diffusion model-based methods, we compare with DPS (Chung et al., 2023), DDNM (Wang et al., 2023), and DiffPIP (Zhu et al., 2023). These methods all use the same pre-trained diffusion models as priors, but each employs a distinct posterior sampling strategy to solve general inverse problems. We specifically use the pre-trained diffusion model from DiffPIP. We also compare our approach with DRP (Hu et al., 2024c), a method that leverages deep restoration priors for solving general inverse problems. For a fair comparison, we use the same pre-trained deblurring network as in ShaRP for DRP. However, instead of employing a set of degradation priors, DRP uses a single fixed prior. For a fair comparison, we selected the optimal fixed prior—defined by a fixed  $\alpha$  based on PSNR performance on the validation set, and applied it accordingly. For all baselines, we fine-tuned their parameters to maximize PSNR performance. Notably, the diffusion model backbone for all diffusion-based baselines was trained on the same dataset used to train our restoration prior.

<sup>4</sup><https://github.com/NVLabs/I2SB>

## C. Additional results for CS-MRI

### C.1. Performance of ShaRP for uniform and random subsampling setting

Due to space constraints, we present only the quantitative performance for the  $4\times$  uniform subsampling setting in the main paper. In this section, we further evaluate ShaRP’s performance on both uniform random subsampling setting, with two sub-sampling rates ( $4\times$  and  $6\times$ ), and three noise levels ( $\sigma = 0.005, 0.01, \text{ and } 0.015$ ).

Noise level	4× Uniform						6× Uniform					
	$\sigma = 0.005$		$\sigma = 0.010$		$\sigma = 0.015$		$\sigma = 0.005$		$\sigma = 0.010$		$\sigma = 0.015$	
Metrics	PSNR	SSIM										
Zero-filled	26.93	0.848	26.92	0.847	26.90	0.848	22.62	0.728	22.60	0.726	22.59	0.721
TV	31.17	0.923	31.08	0.921	30.91	0.915	25.00	0.806	24.94	0.803	24.33	0.755
PnP-FISTA	35.88	0.938	31.14	0.894	30.32	0.846	26.30	0.822	25.97	0.786	25.46	0.747
PnP-ADMM	<u>35.76</u>	<u>0.941</u>	32.36	0.878	30.66	0.838	26.13	0.808	25.90	0.776	25.51	0.742
DRP	35.52	0.936	32.32	0.914	30.57	0.901	29.51	0.872	28.52	0.882	28.35	0.876
DPS	32.62	0.888	31.39	0.870	30.29	0.856	30.53	0.862	29.41	0.843	28.63	0.830
DDS	35.21	0.937	<u>35.03</u>	<u>0.935</u>	<u>34.51</u>	<u>0.925</u>	<u>31.02</u>	<u>0.889</u>	<u>30.84</u>	<u>0.888</u>	<u>30.79</u>	<u>0.888</u>
ShaRP	<b>37.59</b>	<b>0.963</b>	<b>35.81</b>	<b>0.951</b>	<b>34.92</b>	<b>0.942</b>	<b>33.42</b>	<b>0.940</b>	<b>32.86</b>	<b>0.932</b>	<b>32.09</b>	<b>0.922</b>

Table 5: Quantitative comparison of ShaRP with several baselines for CS-MRI using uniform masks at undersampling rates of 4 and 6 on fastMRI dataset. The **best** and second best results are highlighted. Notably, ShaRP outperforms SOTA methods based on denoisers and diffusion models.

Noise level	4× Random						6× Random					
	$\sigma = 0.005$		$\sigma = 0.010$		$\sigma = 0.015$		$\sigma = 0.005$		$\sigma = 0.010$		$\sigma = 0.015$	
Metrics	PSNR	SSIM										
Zero-filled	25.83	0.815	25.81	0.812	25.76	0.807	22.68	0.724	22.67	0.722	22.67	0.719
TV	28.14	0.866	28.06	0.863	27.96	0.859	24.55	0.782	24.33	0.750	24.28	0.736
PnP-FISTA	29.31	0.863	28.40	0.817	27.49	0.799	26.01	0.797	25.63	0.756	24.94	0.717
PnP-ADMM	28.83	0.842	28.39	0.816	27.70	0.786	25.59	0.776	25.19	0.740	24.93	0.728
DRP	29.97	0.880	29.37	0.839	28.31	0.794	26.98	0.866	26.78	0.853	26.49	0.821
DPS	31.72	0.874	30.45	0.857	29.50	0.843	30.32	0.856	29.36	0.824	27.99	0.810
DDS	<u>32.41</u>	<u>0.910</u>	<u>32.37</u>	<u>0.906</u>	<u>32.25</u>	<u>0.901</u>	<u>30.59</u>	<u>0.876</u>	<u>30.35</u>	<u>0.874</u>	<u>30.31</u>	<u>0.879</u>
ShaRP	<b>34.66</b>	<b>0.949</b>	<b>33.57</b>	<b>0.920</b>	<b>33.18</b>	<b>0.931</b>	<b>31.53</b>	<b>0.924</b>	<b>31.46</b>	<b>0.918</b>	<b>31.45</b>	<b>0.914</b>

Table 6: Quantitative comparison of ShaRP with several baselines for CS-MRI using random masks at undersampling rates of 4 and 6 on fastMRI dataset. The **best** and second best results are highlighted. Notably, ShaRP outperforms SOTA methods based on denoisers and diffusion models.

Table 6 provides a quantitative comparison of reconstruction performance across different acceleration factors and noise levels using a uniform sub-sampling mask. In all configurations, ShaRP consistently outperforms the baseline methods. The use of a set of restoration operators clearly enhances ShaRP’s performance, highlighting the effectiveness of employing multiple operators to maximize the regularization information provided by the restoration model. Figure 6 presents visual reconstructions for two test scenarios, where ShaRP accurately recovers fine brain details, particularly in the zoomed-in regions, while baseline methods tend to oversmooth or introduce artifacts. These results highlight ShaRP’s superior ability to manage structured artifacts and preserve fine details, outperforming both denoiser-based and diffusion model-based methods.

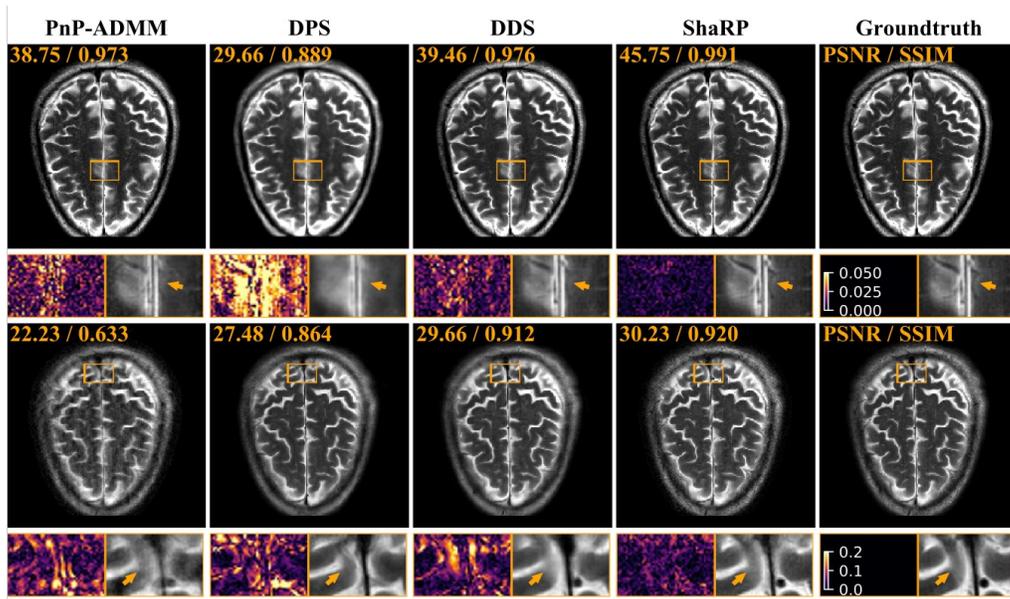


Figure 5: Visual comparison of ShaRP with baseline methods on CS-MRI. The top row shows results for a  $4\times$  random mask with noise  $\sigma = 0.005$ , and the bottom row for a  $6\times$  random mask with noise  $\sigma = 0.015$ . PSNR and SSIM values are in the top-left corner of each image. Error maps and zoomed-in areas highlight differences. Notably, ShaRP with stochastic priors outperforms state-of-the-art methods using denoiser and diffusion model priors.

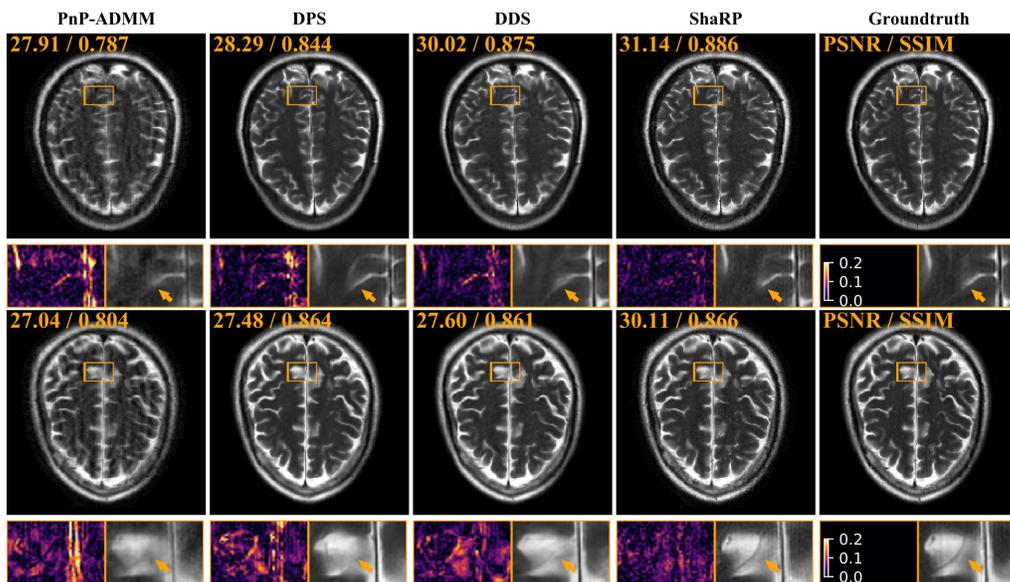


Figure 6: Visual comparison of ShaRP with baseline methods on CS-MRI for  $6\times$  random sampling mask with noise  $\sigma = 0.015$ . PSNR and SSIM values are in the top-left corner of each image. Error maps and zoomed-in areas highlight differences. Notably, ShaRP with stochastic priors outperforms state-of-the-art methods using denoiser and diffusion model priors.

C.2. Performance of ShaRP with Self-supervised restoration priors

Noise level	4× Random						6× Random					
	σ = 0.005		σ = 0.010		σ = 0.015		σ = 0.005		σ = 0.010		σ = 0.015	
Metrics	PSNR	SSIM										
PnP-ADMM	28.83	0.842	28.39	0.816	27.70	0.786	25.59	0.776	25.19	0.740	24.93	0.728
ADMM-TV	28.14	0.866	28.06	0.863	27.96	0.859	24.55	0.782	24.33	0.750	24.28	0.736
GRAPPA	28.09	0.792	25.39	0.699	23.94	0.649	25.67	0.737	23.72	0.646	22.51	0.595
SPICER	<u>31.87</u>	<u>0.901</u>	<u>31.67</u>	<u>0.889</u>	<u>31.50</u>	<u>0.887</u>	<u>30.18</u>	<u>0.871</u>	<u>30.05</u>	<u>0.863</u>	<u>30.01</u>	<u>0.860</u>
ShaRP <sup>self</sup>	<b>33.87</b>	<b>0.909</b>	<b>33.64</b>	<b>0.900</b>	<b>33.21</b>	<b>0.892</b>	<b>30.87</b>	<b>0.899</b>	<b>30.36</b>	<b>0.890</b>	<b>30.21</b>	<b>0.875</b>

Table 7: PSNR (dB) and SSIM values for ShaRP with a self-supervised pre-trained restoration operator, compared to several baselines for CS-MRI with random undersampling masks at rates of 4 and 6 on the fastMRI dataset. The **best** and **second best** results are highlighted. For reference, the highlighted row presents a PnP method using a Gaussian denoiser, which requires fully sampled data for training. Note the excellent performance of ShaRP even using priors trained without fully-sampled ground-truth data.

C.3. Convergence Performance of ShaRP with supervised and self-supervised priors

Figure 7 illustrates the convergence behavior of ShaRP with both supervised and self-supervised learned restoration priors on the test set with an acceleration factor of  $R = 6$  and additional noise  $\sigma = 0.01$ . As observed, ShaRP converges stably under both settings. However, the self-supervised prior exhibits a performance gap compared to the supervised one, indicating that it does not serve as a perfect MMSE estimator in practice. Despite this, it still enables stable convergence as an ensemble of priors for ShaRP, which aligns with our convergence analysis for inexact MMSE estimators presented in Proposition 1.

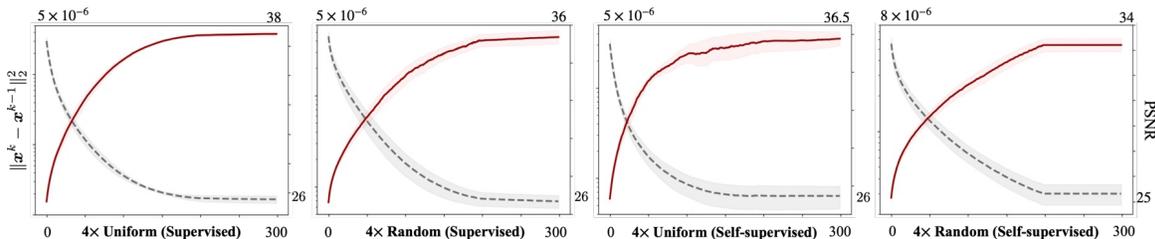


Figure 7: Convergence of ShaRP for 4× accelerated MRI reconstruction on the fastMRI dataset. (a)-(b) depict the convergence behavior of ShaRP using restoration operators trained in a supervised manner, while (c)-(d) correspond to those trained in a self-supervised manner. The plots illustrate the average distance  $\|\mathbf{x}^k - \mathbf{x}^{k-1}\|_2$  and PSNR relative to the ground truth, as a function of the iteration number, with shaded regions representing the standard deviation. Note the stable convergence of ShaRP with both types of priors.

#### C.4. Performance of additional baseline methods on matched and mismatched settings

In this section, we highlight an important observation: pre-trained restoration networks typically exhibit poor generalization to mismatched settings. We chose two commonly used methods (SwinIR (Liang et al., 2021) and E2E-VarNet (Sriram et al., 2020)) for the specific setting of CS-MRI. We trained them on the same  $8\times$  uniform subsampling setting as our restoration prior and directly applied them to solve both matched and mismatched problems, as ShaRP did. As shown in the Table 8, the baseline method’s performance drops significantly under mismatched conditions, whereas ShaRP maintains stable performance and convergence guarantees. This demonstrates ShaRP’s ability to adapt pre-trained restoration models as priors and use it to solve problems under mismatched settings. As shown in the Figure 8, due to the mismatched settings, the two baseline methods suffer from over-smoothing, lack important details, and exhibit artifacts, whereas ShaRP still provides high-quality reconstruction performance. This indicates that ShaRP can balance data fidelity and the artifact removal capabilities of the prior model, leading to an artifact-free reconstruction that preserves important details.

Settings	4× Uniform		4× Random		6× Uniform		6× Random		8× Uniform	
Metrics	PSNR	SSIM								
SwinIR	24.78	0.849	25.09	0.841	29.55	0.907	27.98	0.819	29.37	0.898
E2E-VarNet	<u>35.40</u>	<u>0.957</u>	<u>33.48</u>	<u>0.945</u>	<u>32.79</u>	<u>0.936</u>	<u>31.02</u>	<u>0.913</u>	<b>32.59</b>	<b>0.919</b>
ShaRP	<b>37.59</b>	<b>0.963</b>	<b>34.66</b>	<b>0.949</b>	<b>33.42</b>	<b>0.940</b>	<b>31.53</b>	<b>0.924</b>	<u>32.37</u>	<u>0.907</u>

Table 8: Quantitative comparison of ShaRP with task-specific baselines trained on the  $8\times$  uniform mask. Baselines perform well in matched settings (highlighted in the table) but show a significant drop under mismatched conditions. In contrast, ShaRP remains robust, handling both matched and mismatched scenarios effectively.

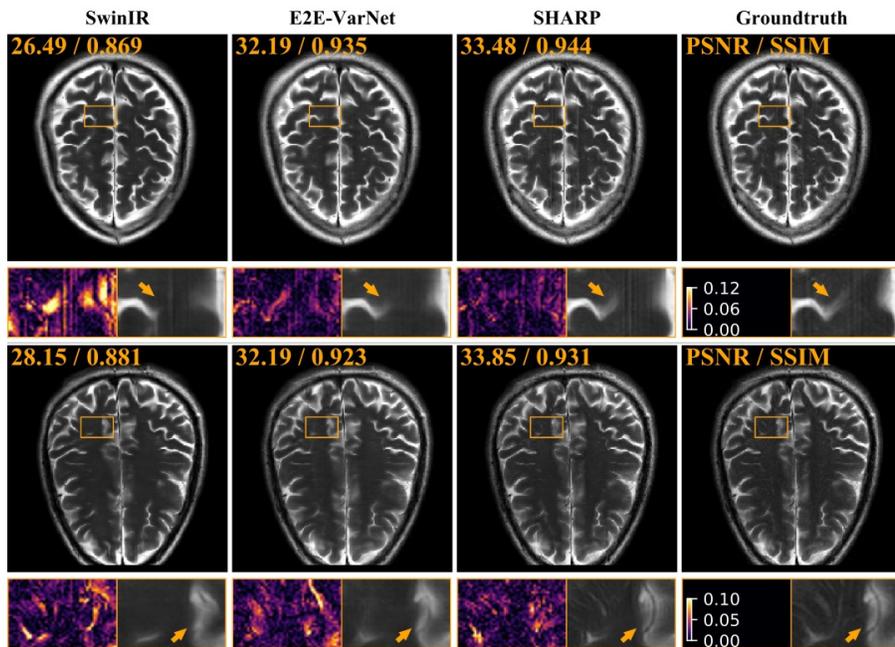


Figure 8: Visual comparison of ShaRP with task-specific baseline methods on CS-MRI for  $6\times$  random sampling mask with noise  $\sigma = 0.015$ . PSNR and SSIM values are in the top-left corner of each image. Error maps and zoomed-in areas highlight differences. Notably, ShaRP with stochastic priors outperforms state-of-the-art methods using denoiser and diffusion model priors.

#### C.5. Additional Evaluation: Non-Cartesian CS-MRI Generalization

In this section, we present an additional evaluation designed to further assess ShaRP’s generalization capabilities. This new experimental setting explores non-Cartesian sampling for Compressed Sensing MRI (CS-MRI), specifically employing a 2D

Gaussian sampling mask. A key aspect of this evaluation is that the restoration network, originally trained for  $8\times$  uniform Cartesian sampling, was directly applied to this unseen sampling scheme without any retraining or fine-tuning. The results, summarized in Table 9, provide compelling evidence of ShaRP’s robustness to diverse sampling patterns and highlight its strong generalization performance.

Table 9: Performance comparison on CS-MRI with a 2D Gaussian sampling mask. The restoration network was trained on  $8\times$  uniform Cartesian sampling and applied directly without retraining or adaptation.

<b>Method</b>	PSNR	SSIM
PnP-ADMM	31.33	0.917
DPS	32.01	0.904
DDS	33.19	0.924
DRP	32.70	0.921
ShaRP	<b>34.01</b>	<b>0.942</b>

### D. Additional visual results for SISR

In this section, we present additional visual results to numerical comparisons for the SISR task.

#### D.1. Additional visual results against baselines

As illustrated in Figure 9 and Figure 10, ShaRP outperforms all baseline approaches under both blur kernel settings, achieving higher PSNR and SSIM values. Moreover, we maintain superior data consistency with the measurements while achieving enhanced perceptual quality. The use of an ensemble of deblurring priors enables our method to recover fine details at varying corruption levels, contributing to the improved performance.

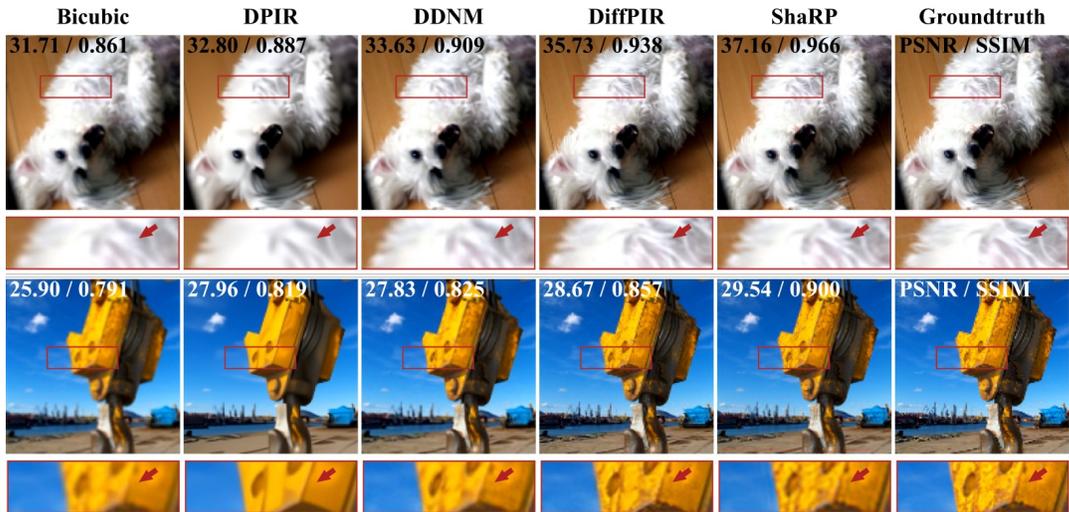


Figure 9: Visual comparison of ShaRP with several well-known methods on  $2\times$  SISR with gaussian blur kernel with  $\sigma = 1.5$ . The quantities in the top-left corner of each image provide PSNR and SSIM values for each method. The squares at the bottom of each image visualize the zoomed area in the image.

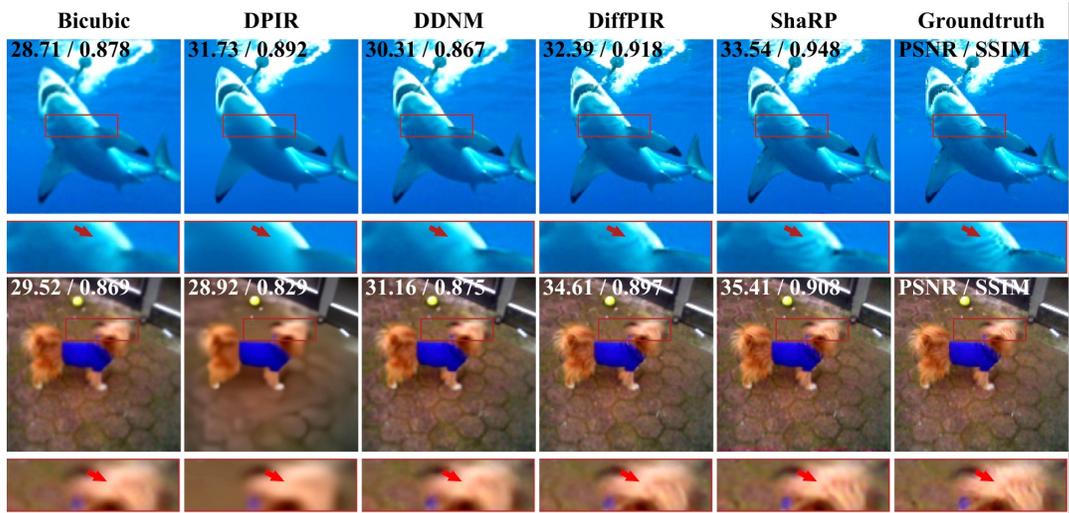


Figure 10: Visual comparison of ShaRP with several well-known methods on  $2\times$  SISR with gaussian blur kernel with  $\sigma = 1.25$ . The quantities in the top-left corner of each image provide PSNR and SSIM values for each method. The squares at the bottom of each image visualize the zoomed area in the image.

## D.2. Additional visual results against DRP

To further emphasize the necessity and advantages of using an ensemble of deblurring priors, as opposed to a fixed prior like in DRP (Hu et al., 2024c), we provide additional visual comparison results. As shown in Figure 11, ShaRP consistently recovers finer details, resulting in improved PSNR and SSIM scores, along with enhanced perceptual performance.

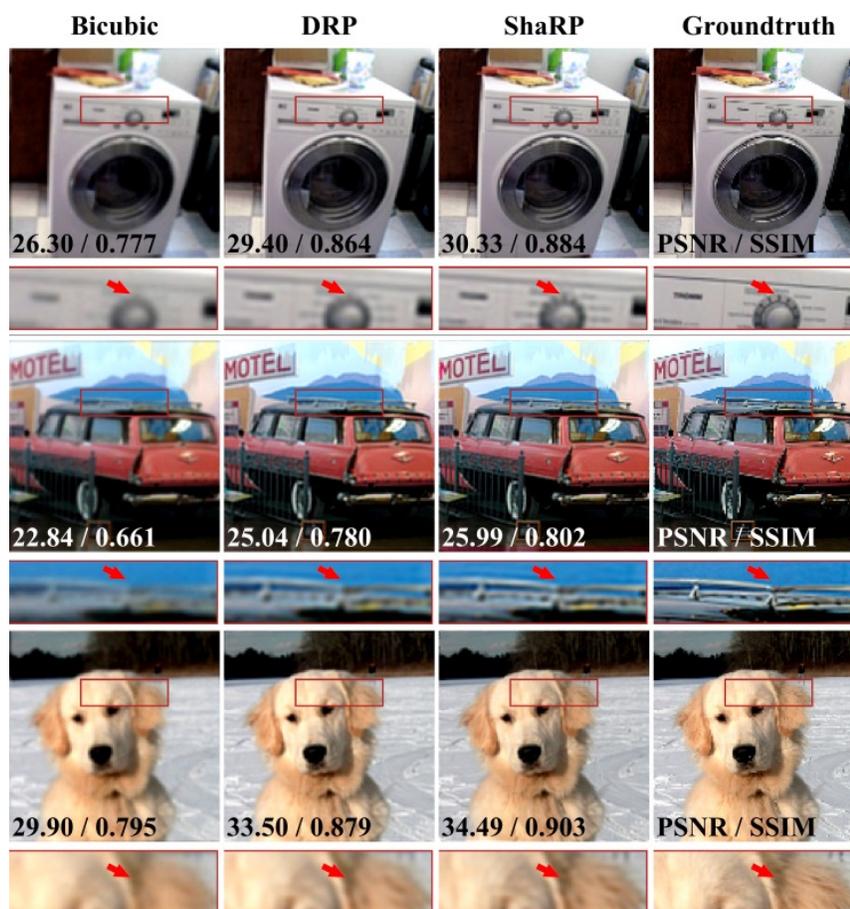


Figure 11: Visual comparison of ShaRP with DRP on  $2\times$  SISR with gaussian blur kernel with  $\sigma = 1.5$ . The quantities in the bottom-left corner of each image provide PSNR and SSIM values for each method. The squares at the bottom of each image visualize the zoomed area in the image.

### D.3. Additional Comparison against DDRM and DiffIR

To further evaluate ShaRP’s performance against state-of-the-art diffusion-based methods, we included two additional baselines for comparison: DDRM (Kawar et al., 2022) and DiffIR (Xia et al., 2023). The experiment setting is  $2\times$  SISR task with gaussian blur kernel with  $\sigma = 1.25$  on ImageNet dataset. For DDRM, we utilized the same pre-trained unconditional diffusion backbone as DiffPIR, DDNM, and DDS, but followed the sampling procedure outlined in their original paper. For DiffIR, we directly used the provided checkpoint from the authors.

Metrics	PSNR	SSIM	LPIPS
DPIR	28.10	0.809	0.305
DDNM	27.53	0.786	0.240
DPS	24.68	0.661	0.395
DiffPIR	28.92	0.852	<b>0.152</b>
DiffIR	25.79	0.812	0.180
DDRM	28.20	0.845	<u>0.161</u>
DRP	<u>29.28</u>	<u>0.868</u>	0.207
ShaRP	<b>30.09</b>	<b>0.891</b>	0.179

Table 10: Quantitative comparison of ShaRP with several additional baselines for  $2\times$  SISR with gaussian blur kernel with  $\sigma = 1.25$  on ImageNet dataset. The **best** and second best results are highlighted. Notably, ShaRP outperforms SOTA methods based on denoisers and diffusion models.

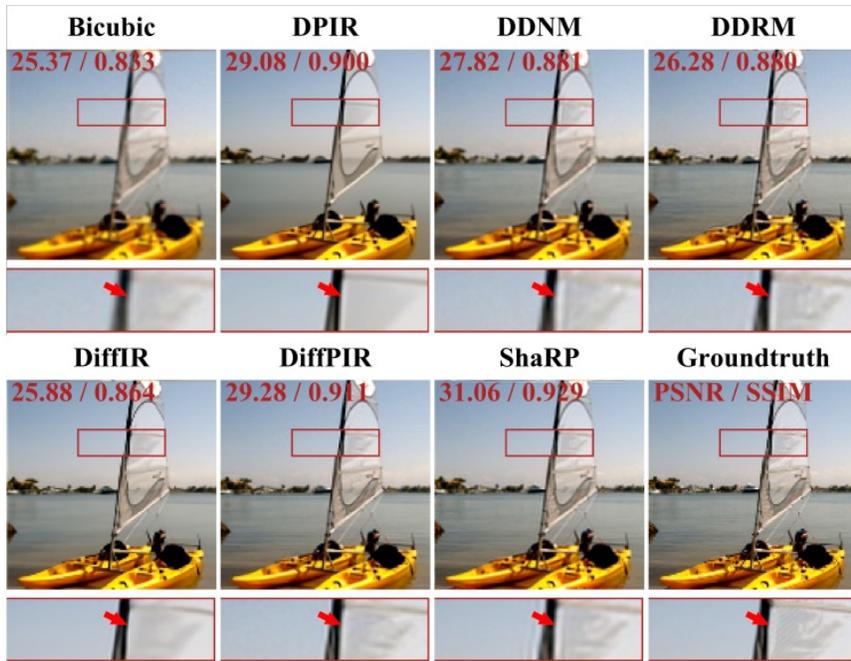


Figure 12: Visual comparison of ShaRP with additional baselines on  $2\times$  SISR with gaussian blur kernel with  $\sigma = 1.25$ . The quantities in the bottom-left corner of each image provide PSNR and SSIM values for each method. The squares at the bottom of each image visualize the zoomed area in the image.

**D.4. Additional Evaluation: SISR with New Gaussian Kernel**

To further investigate ShaRP’s robustness, we ran an additional Super-Resolution (SISR) experiment using a new Gaussian blur kernel with  $\sigma = 1.0$ . This test aimed to evaluate if ShaRP can maintain consistently good performance, comparable to its results with the other two kernels detailed in the main manuscript, even with this previously unseen degradation. As shown in Table 11, ShaRP continues to provide excellent performance, suggesting its strong generalization capabilities and potential effectiveness across a wider range of Gaussian blur levels.

Table 11: SISR performance with a new Gaussian blur kernel ( $\sigma = 1.0$ ).

<b>Method</b>	PSNR (dB)	SSIM	LPIPS ↓	FID ↓
DPIR	28.45	0.854	0.247	82.90
DDRM	27.26	0.803	0.209	44.77
DiffPIR	28.37	0.841	0.215	40.59
DRP	28.43	0.853	0.236	75.29
ShaRP	<b>28.70</b>	<b>0.858</b>	<b>0.226</b>	<b>69.75</b>

## E. Additional Experiments

In this section, we include two additional ablation studies to further highlights ShaRP’s capability to leverage restoration priors for solving general inverse problems, as well as to evaluate its performance under different hyperparameter settings.

### E.1. Ablation study on using SR prior for CS-MRI task

To demonstrate the flexibility of our approach in integrating diverse restoration models to address general inverse problems, we conducted an additional ablation study using a pre-trained super-resolution network as a prior for solving the CS-MRI problem.

---

#### Algorithm 5 MRI Super Resolution network training

---

**Require:** dataset:  $p(\mathbf{x}, \mathbf{y})$ ,  $4\times$  bicubic downsampling operator:  $\mathbf{K}$ ,  $\mathbf{R}_\theta(\cdot, \alpha)$

**repeat:**

$$\mathbf{x} \sim p(\mathbf{x}), \mathbf{e} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}), \alpha \sim \mathcal{U}([0, 1])$$

$$\min_{\theta} \|\mathbf{R}_\theta((1 - \alpha)\mathbf{x} + \alpha \mathbf{D}^\top \mathbf{D} \mathbf{x}; \alpha) - \mathbf{x}\|_2^2$$

**until converge**

---

**Models training for MRI-SR** We use the same U-Net architecture as employed in the official implementation of DDS<sup>5</sup> for  $\mathbf{R}(\cdot; \alpha)$ . To create an ensemble of restoration priors, we consider a family of degradation operators that are convex combinations of the identity mapping  $\mathbf{I}$  and the Gaussian blur operator  $\mathbf{D}$ . The  $4\times$  bicubic downsampling operator  $\mathbf{D}$  corresponds to bicubic downsample with factor equals to 4. Specifically, we define the degradation operator as  $\mathbf{H}_\alpha = (1 - \alpha)\mathbf{I} + \alpha \mathbf{D}^\top \mathbf{D}$ , where  $\alpha \in [0, 1]$  controls the degradation level. By varying  $\alpha$ , we generate multiple degradation operators, allowing us to train the restoration network  $\mathbf{R}$  to handle all these operators, expressed as  $\mathbf{R}(\mathbf{s}, \mathbf{H}_\alpha) = \mathbb{E}[\mathbf{x}|\mathbf{s}, \mathbf{H}_\alpha]$ , where  $\mathbf{s}$  is the degraded image and  $\mathbf{x}$  is the original image.

We select 1,000 different  $\alpha$  values from the interval  $[0, 1]$ , following the  $\alpha$  schedule outlined by I<sup>2</sup>SB (Liu et al., 2023). This results in 1,000 different degradation operators  $\mathbf{H}_\alpha$ , effectively creating an ensemble of restoration priors during training. The model is trained using the Adam optimizer with a learning rate of  $5 \times 10^{-5}$ .

**Using MRI-SR model as prior for CS-MRI task.** During inference, to simplify computation and focus on the most effective priors, we use only a subset of the ensemble. Specifically, we select 6  $\alpha$  values, resulting in 6 restoration priors.

As shown in Table 12, under the  $4\times$  uniform mask setting, employing the pre-trained MRI-SR model as prior allows ShaRP to outperform denoiser- and diffusion-based approaches. However, its performance remains inferior to ShaRP with a mismatched CS-MRI-specific prior. In the  $4\times$  random mask setting, ShaRP with the pre-trained MRI-SR model as prior continues to surpass PnP-based methods that utilize a denoiser prior but performs worse than approaches based on diffusion models. Notably, ShaRP with a mismatched CS-MRI-specific prior consistently delivers the best performance.

Tasks	Metrics	PnP-FISTA	PnP-ADMM	DPS	DDS	ShaR <sub>CS</sub>	ShaR <sub>SR</sub>
4x Uniform	PSNR	35.88	35.76	32.62	35.21	<b>37.59</b>	<u>35.91</u>
	SSIM	0.938	0.941	0.888	0.937	<b>0.961</b>	<u>0.943</u>
4x Random	PSNR	29.31	28.83	31.72	<u>32.41</u>	<b>34.66</b>	30.91
	SSIM	0.863	0.842	0.874	<u>0.910</u>	<b>0.949</b>	0.905

Table 12: Quantitative comparison of ShaRP against baselines for CS-MRI reconstruction using  $8\times$  CS-MRI and  $4\times$  super-resolution priors, evaluated on the fastMRI dataset. Results are reported for both uniform and random undersampling masks at a 4x undersampling rate. The **best** and second best results are highlighted.

<sup>5</sup><https://github.com/HJ-harry/DDS>

**E.2. Ablation Study on the number of restoration priors included in the ensemble**

To evaluate the impact of the number of restoration priors,  $b$ , on ShaRP’s performance, we conducted an ablation study in the  $4\times$  CS-MRI setting with random sampling masks. We specifically examined how varying  $b$  influenced reconstruction performance, offering valuable insights into ShaRP’s sensitivity to this parameter and its role in achieving optimal results.

As illustrated in Figure 13 and Figure 14, increasing  $b$ , which corresponds to incorporating more restoration priors in the ensemble, generally enhances ShaRP’s reconstruction performance.

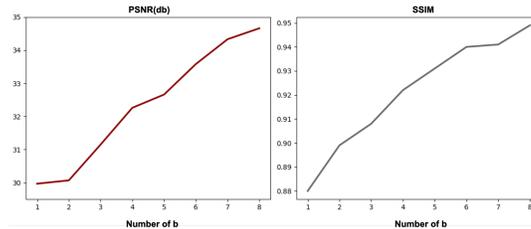


Figure 13: Performance comparison of ShaRP’s CS-MRI reconstruction at  $4\times$  acceleration with varying numbers of restoration priors,  $b$ . Left: PSNR vs.  $b$ ; Right: SSIM vs.  $b$ . ShaRP with  $b = 8$  consistently achieves superior results, highlighting the performance improvements gained by incorporating more restoration priors into ShaRP.

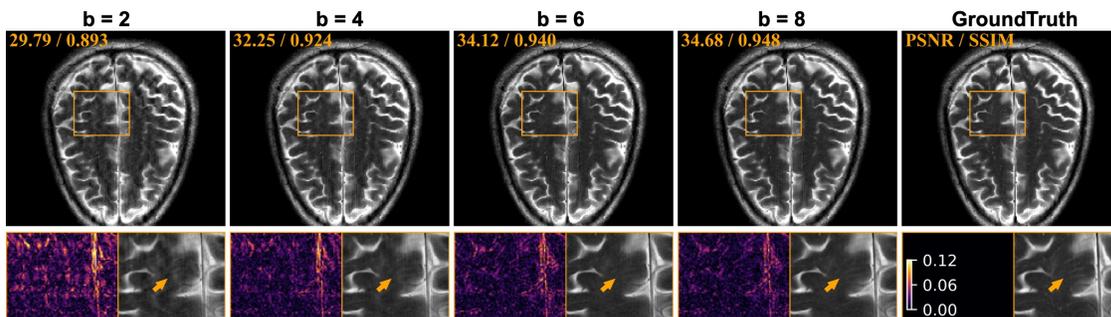


Figure 14: Visual comparison of ShaRP with varying amounts of restoration priors, denoted by  $b$ , in the ensemble. The PSNR and SSIM values for each method are shown in the top-left corner of each image. Zoomed-in regions, highlighted as squares at the bottom of each image, provide a closer look at key details. Notably, increasing the number of restoration priors in the ensemble enhances visual performance by effectively reducing artifacts and capturing finer details.

### E.3. Ablation Study on the Influence of Hyperparameter $\alpha$

To evaluate the impact of the hyperparameters  $\alpha$  as introduced in Section B.1, we conducted an ablation study in the  $4\times$  CS-MRI setting using random sampling masks, where  $\alpha$  governs the selection of a specific restoration prior. Specifically, we analyzed how varying the values of  $\alpha$  influenced reconstruction performance. As shown in Figure 15 demonstrates the influence of  $\alpha$  on performance. A very small  $\alpha$  fails to provide sufficient regularization to constrain the solution, while an excessively large  $\alpha$  overly restricts the model, leading to a decline in performance. These findings highlight the importance of appropriately tuning  $\alpha$  and  $b$  to balance flexibility and regularization for optimal results.

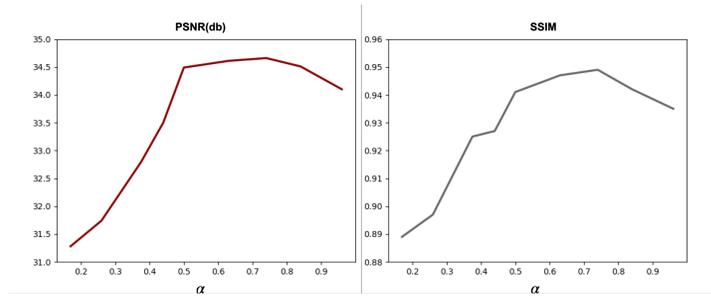


Figure 15: Performance comparison of ShaRP’s CS-MRI reconstruction at  $4\times$  acceleration with varying  $\alpha$ . Left: PSNR vs.  $\alpha$ ; Right: SSIM vs.  $\alpha$ .