A Self-supervised Diffusion Bridge for MRI Reconstruction

Harry Gao*, Weijie Gan*, Yuyang Hu, Hongyu An, and Ulugbek S. Kamilov

Washington University in St. Louis, MO, USA

*These authors contributed equally.

{harrygao, weijie.gan, h.yuyang, hongyuan, kamilov}@wustl.edu

Abstract

Diffusion bridges (DBs) are a class of diffusion models that enable faster sampling by interpolating between two paired image distributions. Training traditional DBs for image reconstruction requires high-quality reference images, which limits their applicability to settings where such references are unavailable. We propose *SelfDB* as a novel *self-supervised* method for training DBs directly on available noisy measurements without any high-quality reference images. SelfDB formulates the diffusion process by further sub-sampling the available measurements two additional times and training a neural network to reverse the corresponding degradation process by using the available measurements as the training targets. We validate SelfDB on compressed sensing MRI, showing its superior performance compared to the denoising diffusion models.

1   Introduction

Many problems in biomedical imaging can be viewed as inverse problems, where the goal is to reconstruct an unknown image from its noisy and incomplete measurements. For instance, in compressed sensing MRI (CS-MRI) [1], the goal is to reconstruct high-quality images from undersampled k-space measurements. Numerous algorithms have been proposed to address inverse problems over the past years with the current research focused on deep learning (DL)-based methods [2,3]. Among DL approaches, diffusion models have gained popularity for their ability to capture complex data distributions and produce images with impressive perceptual quality [4,5].

Denoising diffusion models (DDMs) are a class of methods that define a diffusion process from the ground truth image distribution to the standard Gaussian distribution [6]. A deep neural network (DNN) is then trained to reverse this process, allowing generation of images from randomly generated Gaussian noise. Recent studies have shown the potential of DDMs on the inverse problems via posterior sampling [7] or conditional generation [8]. However, DDMs are known to be slow during inference due to many inference steps. Diffusion bridges (DBs) are alternatives to DDMs based on a diffusion process that bridges the ground truth image distribution directly to the measurement distribution [9,10]. Once a DNN is trained to reverse the this process, inverse problems can be solved by inferring from the pre-trained DB starting from the unseen measurements. Recent work has shown that DBs can outperform DDMs using much fewer inference steps [9].

Despite the success of DBs, training them requires high-quality ground-truth images, which may not be available in some applications. On the other hand, recent works on self-supervised learning has investigated training DNNs without any ground-truth images [11–13]. In particular, a recent paper [14] has proposed a self-supervised method for training traditional DDMs. To the best of our knowledge, self-supervised learning has not been explored in DBs. We address this gap by proposing a new self-supervised method for training diffusion bridges (SelfDB). SelfDB is based on a new diffusion process that further degrades the available measurement two additional times using two different operators. This allows training a DNN to reverse the SelfDB diffusion process using only the available measurements as targets, eliminating the need for ground-truth images. We validate SelfDB on the task of compressed sensing MRI, showing that SelfDB can achieve superior performance compared to a self-supervised DDM.

arXiv:2501.03430v1 [eess.IV] 6 Jan 2025

2 Background

We consider an inverse problem for recovering an unknown image $\boldsymbol{x}$ from its noisy and incomplete measurement

$$\boldsymbol{y} = \boldsymbol{MAx} + \boldsymbol{e} \ , \tag{1}$$

where $\boldsymbol{A}$ is the measurement matrix, $\boldsymbol{e}$ is the measurement noise, and $\boldsymbol{M}$ is a binary sub-sampling matrix. The matrix $\boldsymbol{A}^{\mathsf{H}}$ denotes a conjugate transpose that maps measurements back to the image domain. In CS-MRI, $\boldsymbol{A}$ can be expressed as $\boldsymbol{A} = \boldsymbol{FS}$, where $\boldsymbol{F}$ is the Fourier transform, and $\boldsymbol{S}$ consists of coil sensitivity maps. Our goal is to address the ill-posed nature of this inverse problem by training a DB using only measurements $\boldsymbol{y}$ without any clean images $\boldsymbol{x}$.

2.1 Denoising Diffusion Models

DDMs consider the diffusion process from the image $\boldsymbol{x}$ at $t = 0$ to standard Gaussian noise $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \boldsymbol{I})$ at $t = T$

$$\boldsymbol{x}_t = \sqrt{\bar{\alpha}_t}\boldsymbol{x} + \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon} \ \ \text{for} \ \ t = 1, \ ... \ T, \tag{2}$$

where $\bar{\alpha}_t = \prod_{i=1}^{t} \alpha_i$, and $\{\alpha_t\}_{t=1}^{T}$ are pre-defined parameters. The goal of DDMs is to train a DNN $\boldsymbol{f_\theta}$ to reverse this process, enabling it to obtain a clean image from a noise vector $\boldsymbol{x}_T$ sampled from $\mathcal{N}(0, \boldsymbol{I})$. Specifically, the DNN is trained to denoise noisy images $\boldsymbol{x}_t$ across all different t

$$\mathbb{E}_{\boldsymbol{x},\boldsymbol{\epsilon},t,\boldsymbol{x}_t}\big[\ \|\boldsymbol{f_\theta}(\boldsymbol{x}_t, t) - \boldsymbol{x}\|_2^2\big] \ . \tag{3}$$

During inference, $\boldsymbol{x}$ is estimated from $\boldsymbol{x}_T$ as [6]

$$\boldsymbol{x}_{t-1} = \mu_{1,t}\boldsymbol{x}_t + \mu_{2,t}\boldsymbol{f_\theta}(\boldsymbol{x}_t, t) + \Sigma_t\boldsymbol{\epsilon} \tag{4}$$

where $\mu_{1,t} = \sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})/(1 - \bar{\alpha}_t)$, $\mu_{2,t} = \sqrt{\bar{\alpha}_{t-1}}(1 - \alpha_t)/(1 - \bar{\alpha}_t)$ and $\Sigma_t = (1 - \alpha_t)(1 - \bar{\alpha}_{t-1})/(1 - \bar{\alpha}_t)$. The core concept of this reverse process is to iteratively reduce noise in the image until a clean image is obtained. This framework is designed for *unconditional* image generation. In inverse problems, one can either integrate a measurement consistency term into the reverse iteration (4) for posterior sampling [7] or include the measurements as additional inputs to the network during training (denoted as $\boldsymbol{f_\theta}(\boldsymbol{x}_t, t, \boldsymbol{y})$) [8, 15].

A limitation of traditional DDMs is that they rely on ground-truth images as targets for training. To address this, a recent work on *ambient diffusion* [14] uses a new operator $\boldsymbol{M}'\boldsymbol{A}$ to degrade $\boldsymbol{x}_t$ such that the resulting diffusion process $\boldsymbol{y}_t$ depends only on the measurements $\boldsymbol{y}$

$$\begin{aligned}\boldsymbol{y}_t &= \boldsymbol{M}'\boldsymbol{A}(\sqrt{\bar{\alpha}_t}\boldsymbol{x} + \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon}) \\ &= \sqrt{\bar{\alpha}_t}\boldsymbol{M}'\boldsymbol{y} + \sqrt{1 - \bar{\alpha}_t}\boldsymbol{M}'\boldsymbol{A}\boldsymbol{\epsilon} \ ,\end{aligned} \tag{5}$$

where $\boldsymbol{M}'$ is a further sub-sampling matrix of $\boldsymbol{M}$. A DNN is then trained to reconstruct $\boldsymbol{x}$ from $\boldsymbol{y}_t$ for each t using the measurements as training targets

$$\mathbb{E}_{\boldsymbol{y},\boldsymbol{\epsilon},t,\boldsymbol{y}_t}\big[\ \|\boldsymbol{MAf_\theta}(\boldsymbol{y}_t, t) - \boldsymbol{y}\|_2^2\big] \ . \tag{6}$$

The effectiveness of this training loss has been demonstrated in [14] as well as other related recent work on self-supervised learning for the inverse problems [11, 12]. At inference time, the iteration in (4) is applied by first degrading $\boldsymbol{x}_t$ using $\boldsymbol{M}'\boldsymbol{A}$ and then feeding it to the DNN.

2.2 Diffusion Bridges

We consider DBs that formulate the diffusion process from the ground truth image $\boldsymbol{x}$ at $t = 1$ to the measurement $\boldsymbol{y}$ at $t = 0$ through a simple linear combination [10]

$$\boldsymbol{x}_t = (1 - t)\boldsymbol{x} + t\boldsymbol{A}^{\mathsf{H}}\boldsymbol{y} + \sigma_t\boldsymbol{\epsilon} \ \ \text{for} \ \ t \in [0, 1] \ , \tag{7}$$

Algorithm 1 Inference in SelfDB

1: **input:** $\boldsymbol{y}$, $\boldsymbol{f_\theta}$, δ, $\{\sigma_t\}$.
2: $\boldsymbol{y}_1 = \boldsymbol{M'}\boldsymbol{y} + \sigma_1\boldsymbol{\epsilon}$
3: **for** $t = 1$ to 0 with step δ **do**
4: $\quad\boldsymbol{\epsilon} \sim \mathcal{N}(0, \boldsymbol{I})$
5: $\quad\hat{\boldsymbol{x}} = \boldsymbol{f_\theta}(\boldsymbol{y}_t, t)$
6: $\quad$**if** $t > 0$ **then**
7: $\quad\quad\hat{\boldsymbol{y}} = \overline{\boldsymbol{M}}\boldsymbol{A}\hat{\boldsymbol{x}}$
8: $\quad$**else**
9: $\quad\quad$**return** $\hat{\boldsymbol{x}}$
10: $\quad$**end if**
11: $\quad\boldsymbol{y}_{t-\delta} = \frac{\delta}{t}\hat{\boldsymbol{y}} + (1 - \frac{\delta}{t})\boldsymbol{y}_t + (t-\delta)\sqrt{\sigma_{t-\delta}^2 - \sigma_t^2}\boldsymbol{\epsilon}$.
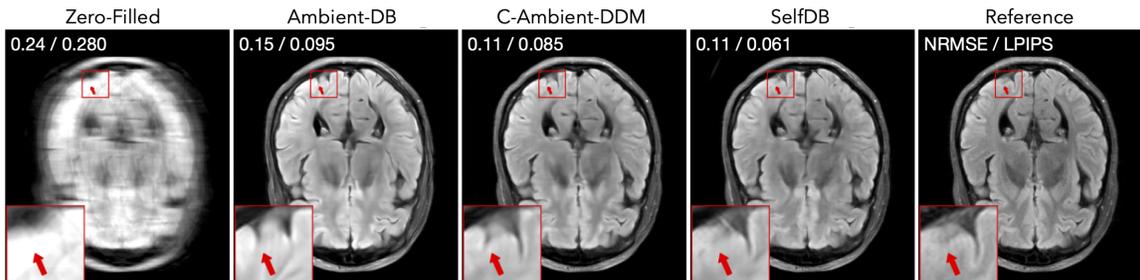12: **end for**

Figure 1: Visual comparisons of *Ambient-DB*, conditional ambient diffusion (*C-Ambient-DDM*), and the proposed SelfDB. *C-Ambient-DDM* is an existing self-supervised approach for denoising diffusion models, while *Ambient-DB* is its direct extension into the DB framework. NRMSE and LPIPS values for each method are labeled in the top-left corner of the images. Note how SelfDB yields images with details closely matching the reference, highlighted by the read arrow in the zoomed-in region.

where σ_t is a standard deviation. Like DDMs, DBs aim to train a neural network to reverse this process, allowing iterative image reconstruction from the measurement $\boldsymbol{y}$ (i.e., $\boldsymbol{x}_t$ at $t = 1$). Because the reverse process of DBs originates from the measurement domain—a closer approximation to the ground truth domain than the standard Gaussian domain in DDM—it offers superior performance with less inference steps [9]. The supervised training loss for DBs can be formulated as

$$\mathbb{E}_{\boldsymbol{x},\boldsymbol{\epsilon},t,\boldsymbol{x}_t}\left[\,\|\boldsymbol{f_\theta}(\boldsymbol{x}_t, t) - \boldsymbol{x}\|_2^2\,\right]\ . \tag{8}$$

During inference, with $\boldsymbol{x}_1 = \boldsymbol{A}^{\mathsf{H}}\boldsymbol{y} + \sigma_1\boldsymbol{\epsilon}$, the DB iteration is given by $\mathbb{E}[\boldsymbol{x}_{t-\delta}|\boldsymbol{x}_t]$ (see also Proposition 4.1 in [10])

$$\boldsymbol{x}_{t-\delta} = \frac{\delta}{t}\boldsymbol{f_\theta}(\boldsymbol{x}_t, t) + (1 - \frac{\delta}{t})\boldsymbol{x}_t + (t-\delta)\sqrt{\sigma_{t-\delta}^2 - \sigma_t^2}\boldsymbol{\epsilon}\ , \tag{9}$$

where δ is the iteration step size. The key concept behind this reverse process is that it gradually improves the image in small steps, making the reconstruction task less and less challenging compared to directly solving the inverse problem.

3 Method

We now introduce our proposed self-supervised algorithm for DBs. We begin by showing why a direct application of ambient diffusion to DBs, which we call *Ambient-DB*, is suboptimal in CS-MRI. Then, we present our proposed SelfDB approach.

3

Table 1: Average NRMSE, SSIM, and LPIPS values for SelfDB, *Ambient-DB*, and *C-Ambient-DDM* on the test set. This table demonstrates that SelfDB outperforms these baselines across both distortion and perceptual metrics.

Metrics	NRMSE ↓	SSIM ↑	LPIPS ↓
Ambient-DB	0.172	0.782	0.069
C-Ambient-DDM	0.113	0.833	0.056
SelfDB	**0.108**	**0.882**	**0.036**

The core concept behind ambient diffusion is to construct a new diffusion process that eliminates the need for ground truth data during training as in (5). A direct adaptation of this concept into DBs is to apply the same operator $\boldsymbol{M}'\boldsymbol{A}$ to $\boldsymbol{x}_t$

$$
\begin{aligned}
\boldsymbol{y}_t &= \boldsymbol{M}'\boldsymbol{A}\boldsymbol{x}_t = \boldsymbol{M}'\boldsymbol{A}((1-t)\boldsymbol{x} + t\boldsymbol{A}^{\mathsf{H}}\boldsymbol{y} + \sigma_t\boldsymbol{\epsilon}) \\
&= \boldsymbol{M}'\boldsymbol{y} + \sigma_t\boldsymbol{M}'\boldsymbol{A}\boldsymbol{\epsilon} ,
\end{aligned}
\tag{10}
$$

where the last line holds because $\boldsymbol{M}'$ is a further sub-sampling matrix of $\boldsymbol{M}$ and $\boldsymbol{A}$ is orthogonal in CS-MRI. Since only a small amount of noise is added during training [10], $\boldsymbol{y}_t$ does not establish a diffusion process that transitions between distinct data, as intended by the concept of DBs. Instead, different values of t correspond to nearly identical sub-sampled data $\boldsymbol{M}'\boldsymbol{y}$, rendering the reverse process effectively as repeated solutions of the same inverse problem. Our experimental results in Section 4 also show that this diffusion process leads to suboptimal performance.

To address this, the proposed SelfDB approach considers a new sub-sampling operator $\overline{\boldsymbol{M}}$ that provides an intermediate level of subsampling between $\boldsymbol{M}$ and $\boldsymbol{M}'$. This enables a diffusion process for SelfDB that transitions from intermediately sub-sampled data at $t = 0$ to highly sub-sampled data at $t = 1$

$$
\boldsymbol{y}_t = (1-t)\overline{\boldsymbol{M}}\boldsymbol{y} + t\boldsymbol{M}'\boldsymbol{y} + \sigma_t\boldsymbol{\epsilon} .
\tag{11}
$$

This modified diffusion process allows for a reverse process that does not directly reconstruct $\boldsymbol{M}'\boldsymbol{y}$. Instead, it incrementally incorporates information from $\overline{\boldsymbol{M}}\boldsymbol{y}$, making the inverse problem progressively less challenging. In alignment with (6), we define our training loss to map $\boldsymbol{y}_t$ to the measurement $\boldsymbol{y}$ as

$$
\mathbb{E}_{\boldsymbol{y},\boldsymbol{\epsilon},t,\boldsymbol{y}_t}\left[\|\boldsymbol{M}\boldsymbol{A}\boldsymbol{f}_{\boldsymbol{\theta}}(\boldsymbol{y}_t,t) - \boldsymbol{y}\|_2^2\right] .
\tag{12}
$$

The inference of SelfDB is summarized in Algorithm 1, where the main update iteration (line 11) is given by $\mathbb{E}[\boldsymbol{y}_{t-\delta}|\boldsymbol{y}_t]$.

4   Numerical Validation

We validated SelfDB in CS-MRI using fully-sampled T2-weighted MR brain data from around 1000 subjects in the fastMRI dataset [16]. These subjects were split into 800 for training, 100 for validation, and 100 for testing. For each subject, we used the middle 7 to 9 slices in the transverse plane. The coil sensitivity map $\boldsymbol{S}$ was estimated via ESPIRiT [17]. For our experiments, we simulated Cartesian sampling operators with sampling rates of 25%, 16.6%, and 12.5% as $\boldsymbol{M}$, $\overline{\boldsymbol{M}}$, and $\boldsymbol{M}'$, respectively. Specifically, we trained SelfDB on a dataset with 25% sub-sampled MRI data and tested it on data sub-sampled at 12.5%. During testing, we considered a small inference step of 4. We evaluated the results using *normalize-root-mean-square-error(NRMSE)/structural similarity index measure (SSIM)* and *learned perceptual image patch similarity (LPIPS)* as distortion based and perceptual metrics, respectively.

We compare SelfDB against *Ambient-DB* and ambient diffusion [14]. Each baseline model was trained using the same experimental setup. Notably, ambient diffusion was originally designed for unconditional
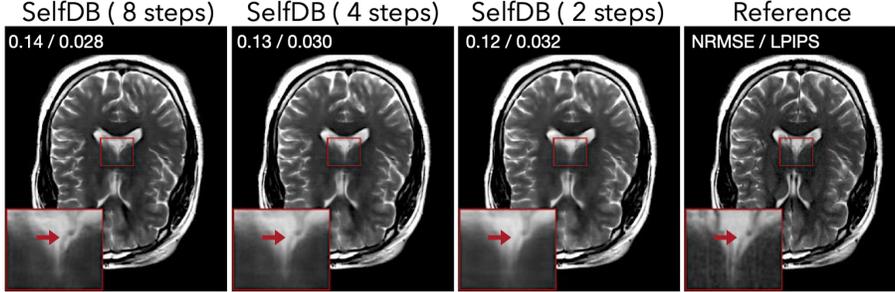
Figure 2: Visual results of SelfDB with different inference steps. Best viewed in digital format. This figure demonstrates that SelfDB empirically provides a perception-distortion trade-off. Notably, perceptual quality (LPIPS) improves as the number of inference steps increases, while distortion with respect to the reference (NRMSE) also rises.

Table 2: Average NRMSE, SSIM, and LPIPS values for SelfDB on the test set across different inference steps. This table shows that the number of inference step in SelfDB empirically controls the perception-distortion trade-off.

Metrics	NRMSE ↓	SSIM ↑	LPIPS ↓
SelfDB (2 steps)	**0.101**	**0.888**	0.039
SelfDB (4 steps)	0.108	0.882	0.036
SelfDB (8 steps)	0.118	0.875	**0.035**

image generation. To adapt it for inverse problems, we implemented it within a conditional diffusion model framework—termed *C-Ambient-DDM*—by incorporating measurements as additional inputs to the DNN.

Figure 1 presents visual results comparing SelfDB with baseline methods, showing that SelfDB can reconstruct images with both high perceptual quality and fine details closely matching the reference. Table 1 summarizes the test metrics for SelfDB and the baseline methods, demonstrating that SelfDB can gain superior performance in both distortion-based and perceptual metrics. In addition to its promising performance, SelfDB also empirically provides a perception-distortion trade-off that can be adjusted by varying the number of inference steps. As illustrated in Figure 2 and shown in Table 2, increasing the sampling steps improves perceptual quality while introducing more distortion. This trade-off is also discussed in other DB-based methods [10]. The ability of SelfDB to balance the perception-distortion trade-off may be particularly valuable in applications where either quantitative accuracy or perceptual quality is preferred.

5 Conclusion

This paper presents SelfDB as the first self-supervised learning algorithm for diffusion bridges. SelfDB proposes a novel diffusion process from intermediately to severely subsampled data. This data is generated by applying two different subsampling operators to the available measurements. The core innovation in this approach lies in *(a)* a diffusion process that bypasses the need for ground truth images and *(b)* allowing the training of a DNN to reverse this process using raw measurements as training targets. We validated SelfDB in CS-MRI. The experimental results show that SelfDB can gain superior performance than its denoising diffusion counterpart under small inference steps.

6   Compliance with Ethical Standards

This research study was conducted retrospectively using human subject data made available in open access by [16]. Ethical approval was not required as confirmed by the license attached with the open access data.

7   Acknowledgments

References

[1] M. Lustig, D. Donoho, and J. M Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, 2007.

[2] G. Ongie et al., "Deep learning techniques for inverse problems in imaging," *IEEE J. Sel. Areas Inf. Theory*, vol. 1, no. 1, pp. 39–56, 2020.

[3] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-Based Deep Learning Architecture for Inverse Problems," *IEEE Transactions on Medical Imaging*, vol. 38, no. 2, pp. 394–405, Feb. 2019.

[4] G. Daras et al., "A Survey on Diffusion Models for Inverse Problems," *arXiv:2410.00083*, 2024.

[5] J. Liu et al., "DOLCE: A model-based probabilistic diffusion framework for limited-angle ct reconstruction," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2023, pp. 10498–10508.

[6] J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, vol. 33, pp. 6840–6851.

[7] H. Chung et al., "Diffusion Posterior Sampling for General Noisy Inverse Problems," in *Proc. Int. Conf. Learn. Represent.*, Feb. 2023.

[8] Y. Mao et al., "DisC-Diff: Disentangled Conditional Diffusion Model for Multi-Contrast MRI Super-Resolution," in *Proc. Med. Image. Comput. Comput. Assist. Intervent*, 2023, pp. 387–397.

[9] H. Chung, J. Kim, and J. C. Ye, "Direct Diffusion Bridge using Data Consistency for Inverse Problems," in *Proc. Adv. Neural Inf. Process. Syst.*, 2023.

[10] M. Delbracio and P. Milanfar, "Inversion by Direct Iteration: An Alternative to Denoising Diffusion for Image Restoration," *Trans. Mach. Learn.*, 2024.

[11] B. Yaman et al., "Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data," *Magn. Reson. Med.*, vol. 84, no. 6, pp. 3172–3191, 2020.

[12] W. Gan et al., "Self-Supervised Deep Equilibrium Models With Theoretical Guarantees and Applications to MRI Reconstruction," *IEEE Trans. Comput. Imaging*, vol. 9, pp. 796–807, 2023.

[13] D. Chen, J. Tachella, and M. E. Davies, "Equivariant Imaging: Learning Beyond the Range Space," in *Proc. IEEE Int. Conf. Comput. Vis.*, Aug. 2021, pp. 4379–4388.

[14] G. Daras et al., "Ambient diffusion: Learning clean distributions from corrupted data," in *Proc. Adv. Neural Inf. Process. Syst.*, 2023.

[15] C. Saharia et al., "Image Super-Resolution via Iterative Refinement," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2022.

[16] F. Knoll et al., "fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning," *Radiol.: Artif. Intell.*, vol. 2, no. 1, pp. e190007, 2020.

[17] M. Uecker et al., "ESPIRiT-an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA," *Magn. Reson. Med.*, vol. 71, no. 3, pp. 990–1001, 2014.